The Kelly Criterion and Learning in Competitive Insurance Underwriting

Frank Schmid and Can Wang

frank.schmid@frankschmid.com ken.wang.can@gmail.com

This version: 30 October 2026
First version: 17 March 2026

Abstract. A rating plan gives an insurer a price for an account. It does not say what the price is for. This paper gives underwriting an objective, the long-run growth of the insurer's surplus. Under the Kelly criterion an insurer pursues it by maximizing the expected logarithm of surplus at the horizon, which maximizes the expected rate at which surplus compounds. The paper asks what an insurer must know to pursue that objective, and how much of it the insurer can learn from its own operations. Applied to one account, the criterion gives a price, the Kelly decision. The decision needs four things: the account's expected loss; the variability of that loss against the insurer's surplus; the value of keeping the account; and the chance that a quote wins it. A rating plan supplies a version of the first. The insurer can learn all four from what it already holds, with what the market makes public. What it holds is its claim files, the loss runs of the accounts it is asked to quote, its own retention, and the broker's report of the quotes it lost. The paper tests this in a simulated market for middle-market general liability. Three conventional insurers, each pricing by a rating bureau's plan, compete there with the insurer the paper tests, in the market's fourth seat. The fourth prices on one of three estimates of the loss: the rating plan's; a learned estimate, on every account; or the learned estimate on new business with the rating plan's on renewals, which the paper calls learning on new business. It prices through the rating plan, the plan route, or by the Kelly decision, the Kelly route. In the main market, the market at the paper's base values, learning the loss adds more growth than anything else the paper measures, and learning an account's claim frequency is worth more than learning the size of its claims. On the plan route, learning the loss pays on new business; on renewals the rating plan's estimate, computed in part from the insurer's own claim file through the experience modification, grows surplus more than the learned estimate. The plan route with learning on new business grows the most of the insurers the paper reports in the fourth seat in that market. In a market with a second insurer that learns, and in one with a severity trend that no insurer projects, the plan route on the learned estimate grows more than the plan route with learning on new business (under the trend, not at the level that allows for the number of tests); in the market of one tenth of the accounts the difference is not resolved. In the main market, what the Kelly decision adds depends on the estimate it prices. On the rating plan's estimate the Kelly route grows more than the plan route. On the learned estimate it trails the plan route for most of the measured years and ends within an equivalence margin of it, set after the difference was measured. With learning on new business it falls short of the plan route. On the learned estimate the Kelly route's difference from the plan route changes sign across the other values of the main market's dimensions at which the Kelly route was run. Of four further markets, the Kelly route on that estimate grows more than the plan route in three, those of smaller insurers (not at the level that allows for the number of tests), of one tenth of the accounts, and with the severity trend; in the fourth, with a second insurer that learns, it ends within the margin, set after that difference was measured.

JEL classification: G22, D83, L11, C63.

This paper was developed with assistance from generative AI tools, including Claude, across the design, analysis, coding, and drafting stages. The authors supervised all stages, verified the results, and take full responsibility for the content.

Contents

Abbreviations and Acronyms ii

Greek Letters ii

Latin Symbols with More Than One Meaning iii

1 Introduction 1

2 Growth as the Objective, and Learning as Its Source 5

3 A Stylized Market 7

4 What Learning the Loss Is Worth 10

5 The Kelly Decision on Each Estimate 13

6 Other Settings and Other Markets 17

7 What Follows for an Insurer, and Conclusion 20

A The Market: Its Assumptions, Its Insurers, and Its Loss Process 24

B The Rating Plans 30

C The Capital Standard 31

D The Kelly Decision and the Estimators 32

E The Kelly Route's Computation, Step by Step 39

F The Registry of Predictions and Outcomes 40

G The Quote Data and the Further Markets 45

Bibliography 48

Abbreviations and Acronyms

ACORD	Association for Cooperative Operations Research and Development
GL	General liability
GLM	Generalized linear model
IRIS	Insurance Regulatory Information System
NAIC	National Association of Insurance Commissioners
RBC	Risk-based capital
s.e.	Standard error

Greek Letters

α	alpha	Permanent part of a trait (§A.6)
β	beta	Discount factor of the lifetime value: at the insurer's own growth of surplus, the return on funds included, or at the return on funds (§D.1.3, Appendix E; Table 1 names the two discounts)
Δ	Delta	Logarithm of the holder's quote over the lowest invited quote (§A.4)
ε	epsilon	Expenses as a share of premium (§A.3)
ζ	zeta	Latent log price of the hold curve, over the price its deviation is taken from (§D.2.2)
η	eta	Yearly shock to the drifting part of a trait (§A.6)
κ	kappa	The capital requirement's scale (Appendix C)
λ	lambda	An account's expected number of claims in a policy year (§A.1); $\hat{\lambda}$, the tracked claim rate (Equation D.5); λ_0, the claims a year at the median size (§A.6)
μ_ζ	mu	Mean of the hold curve's ζ (§D.2.2)
ν	nu	Mean of the permanent traits' sum, in the schedule modification (Appendix B)
ξ	xi	Standard normal draw of a case reserve's error (§A.6)
π	pi	Provision for expenses and profit in a plan's premium (§A.3); the posterior density of the hold curve's parameters (Equation D.6)
ρ	rho	Persistence of the drifting part of a trait (§A.1); the capital requirement's reserve-charge ratio (Appendix C); with subscript f or s, the fitted persistence of an autocovariance (§D.3)
σ	sigma	Spread of the hold curve's ζ (§D.2.2); with subscript c, of a case reserve's error (§A.6)
τ	tau	Mean time from booking to payment, years (§D.1.2)
Σ	Sigma	Covariance of a file's observations (§D.3); the summation sign (Equation D.3, Appendix B)
ϕ	phi	Standard normal density (§D.3)
Φ	Phi	Standard normal distribution function (§D.2.2)
ψ	psi	Complement of credibility in the loss-ratio method's indication (§A.3)
ω	omega	Weight of a quote in the hold curve's likelihood (§D.3)

Latin Symbols with More Than One Meaning

- A An account's capped losses on its file (§A.3); a policy's capped reported losses (Appendix B); with subscript f or s, a fitted constant of an autocovariance (§D.3)
- a Reported losses, in the loss-ratio method's indication (Equation A.2); the loss adjustment expense's share (Equation D.5)
- B With subscript c, a class's expected capped loss of a year per unit of size (Appendix B); with subscript f or s, a fitted coefficient of an autocovariance (§D.3)
- c An account's class (§A.3); as a subscript, a case reserve (σ_c); c_x, the return on the reserve held (§D.1.2); c_f and c_s, covariances of the observations (§D.3)
- d Base of the weights of the experience modification (§A.3); with subscript it, the drifting part of a trait (§A.6)
- e Expected reported losses, in the loss-ratio method's indication (Equation A.2); elsewhere the base of the natural logarithm
- E A class's expected capped losses, E_{rep} their reported part (§A.3); expected reported losses (Appendix B); $\mathbb{E}$, expectation (§D.1.1); E, the further market of smaller insurers (Table G.3)
- f The change of level on the experience, a/e, before credibility (Equation A.2); the GLM insurer's change of level or a fitted factor (Appendix B)
- g An insurer's growth (Equation A.3); $\hat{g}$, its mean yearly log growth of surplus over the lifetime value's estimation window, the return on funds included (§D.1.3)
- j The age of a year, in the experience modification (§A.3); a quote, in the hold curve's likelihood (§D.3)
- k The buyer's cost of shopping, k' its cost of switching (§A.4); years ahead (§D.1.3); a lag (§D.3)
- L With subscript c, the bureau's loss cost (Equation A.1); L, the further market with a second insurer that learns (Table G.3)
- m With subscript 1 or s, a margin per unit of net premium (§D.1.3); m_0, the mean of a log claim on the bureau's curve (§D.3); without subscript, the equivalence margin (Table F.3)
- n Premium net of expenses (§D.1.2); reported claims (Equation A.2, §D.3), n_s^{open} the open claims among them (§D.3)
- s The scale of the buyer's response (§A.4); as a subscript, later renewals (§D.1.3), the severity trait, or an account-year (§D.3); s_1^2 and s_2^2, variances (Appendix B)
- T The number of measured years (Equation A.3); T, the further market with the severity trend (Table G.3)
- t A year (§D.1.3, §A.6); the t statistic of a difference (Section 1)
- u The occurrence limit (Equation D.5); with subscript it, the frequency trait (§A.6)
- v $\hat{v}$, the estimated variance of an account's loss, $\hat{v}_0$ before the rate level (§D.1.2, §D.2.1); v_0, the variance of a log claim on the bureau's curve (§D.3)
- w With subscript j, the weight of a year in the experience modification (Appendix B); a file's length in years (§D.3)
- x $\bar{x}$, the plan's expected loss (Equation A.1); $\hat{x}$, the estimated expected loss, $\hat{x}_0$ before the rate level (§D.1.2, §D.2.1)
- Y Claim size at the tracked scale (Equation D.5); $\bar{y}$, a mean log claim size (§D.3)
- Z Classical credibility (Equation A.2, Appendix B); Z_{exp}, the experience modification's credibility (§A.3)

1 Introduction

An insurer that competes for an account decides two things at once: at what price to quote, and whether it wants the account at the price the market will bear. It lacks what it needs for both. The first requires the account's loss distribution, which no one observes; the second requires the chance of winning at each price, which depends on what competitors quote and on how the buyer chooses. The industry's answer is the rating plan: a class loss cost, an experience modification, a schedule modification, and a provision for expenses and profit, with the insurer's *rate level*, the multiplier it applies to the class loss costs, moved once a year by what its book has earned against what it was priced to earn.[1] The plan states a price. It does not state what the price is for.

This paper gives underwriting an objective and asks what an insurer needs in order to pursue it. The objective is the long-run growth of surplus. Surplus compounds: what an insurer writes this year changes the surplus from which it writes next year. The Kelly criterion is the rule for this objective. An insurer that maximizes the expected logarithm of its surplus at a horizon maximizes the expected mean geometric rate at which its surplus grows. The horizon is the year at which growth is read, and the logarithm accounts for the compounding. In repeated independent bets on known odds, no other rule grows wealth faster in the long run.[2] The paper uses the criterion in two ways. It is the standard: every insurer in the paper is judged by the growth of its surplus. And it gives a rule for pricing: applied to one account, the criterion gives a price, the *Kelly decision*, from four things: the account's expected loss; the variability of that loss, weighed against the insurer's surplus; the value of keeping the account in the years after this one; and the chance that the quote wins. A rating plan supplies a version of the first and none of the other three.

An insurer can learn all four from what it already holds, with what the market makes public (§D.2). Its claim files record every claim of every account it has written; the loss runs of the accounts it is asked to quote record their claims elsewhere; its own retention of accounts records which stayed at which price; and, where an account went to market, the broker reports by how much a lost quote lost (in the model, rounded; §A.4). From these the insurer learns, and the paper's learning has two parts. *Learning the loss* is the tracking of each account's claim frequency and claim size, updated each year as its files grow. *Learning the hold curve* is fitting the insurer's *hold curve*, its *chance of winning*: the chance that a quote wins the account at each price, on new business the chance of winning an account that goes to market.

The paper measures what this learning is worth in a stylized market for middle-market general liability (GL), which it simulates. Four insurers compete in it for 4,000 accounts, and the paper calls their places in the market *seats*. Three seats hold the same three conventional insurers in every run, except in one of the further markets of Section 6, where a second insurer that learns takes the seat of one of them. Each prices by the rating plan of a rating bureau, which publishes a loss cost for each class of account to every insurer, and each is named by how it sets its rate level. The *Bureau insurer* refiles its rate level when its experience departs from the plan. The *Loss-ratio insurer* moves its rate level each year by the loss-ratio method. The third fits a generalized linear model (GLM) to its own book and adds its relativities by class and size to the plan; it is the *GLM*

[1] Geoff Werner and Claudine Modlin, *Basic Ratemaking*, 5th ed. ([Arlington, VA]: Casualty Actuarial Society, 2016), chaps. 7–9, 12, and 15; Ginda Kaplan Fisher, Lawrence McTaggart, Jill Petker, and Rebecca Pettingell, *Individual Risk Rating* (Casualty Actuarial Society study note, 2019).

[2] John L. Kelly Jr., "A New Interpretation of Information Rate," *Bell System Technical Journal* 35, no. 4 (1956): 917–926; Leo Breiman, "Optimal Gambling Systems for Favorable Games," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, ed. Jerzy Neyman (Berkeley: University of California Press, 1961), 65–78.

insurer. The fourth seat holds the insurer the paper tests. The market is run with each insurer tested in the fourth seat in turn, and nothing else changes: the same three competitors, the same rules, and the same random draws. A difference between two insurers in the fourth seat therefore comes from the insurer in the fourth seat and from how the rest of the market responds to it.

The insurers in the fourth seat differ in two choices. The first is the estimate of each account's loss that the insurer prices on: the rating plan's estimate; the *learned estimate*, the insurer's estimate from the data it holds, on every account; or the learned estimate on new business with the rating plan's estimate on renewals, which the paper calls *learning on new business*. The second is the rule that turns the estimate into a quote: the rating plan, the *plan's pricing*; or the Kelly decision, which moves the price the rating plan gives the estimate, the *undeviated price*, up or down. An insurer that prices its estimate by the plan's pricing is the *plan route* on that estimate, and one that prices it by the Kelly decision the *Kelly route*. The two choices give six insurers in the fourth seat, the cells of a grid (Table 3). Three are priced through the rating plan: the plan route on the rating plan's estimate, a second Loss-ratio insurer, which is the *benchmark*; the plan route on the learned estimate; and the plan route with learning on new business. Three are priced by the Kelly decision: the Kelly route on the rating plan's estimate; the Kelly route on the learned estimate; and the Kelly route with learning on new business. Table 1 adds variations of the Kelly route, 14 insurers in all. The benchmark is a second Loss-ratio insurer because the plan route prices as the Loss-ratio insurer does and differs from it in the estimate alone, so that the difference between the two is what learning the loss is worth at the plan's pricing. The benchmark learns neither the loss nor the hold curve. Whichever estimate of the loss it prices, the Kelly route learns the hold curve from its own quotes won and lost and reads the value of keeping the account from its own book (§2.2); on the rating plan's estimate it does not learn the loss.

Each run of the market opens with a build-up of the books and 10 practice years, so that every insurer's estimators have run before the 28 *measured years*, over which its growth is read. The market is run 200 times. Each run is one draw of the market's random events, the accounts' claims among them, and the paper calls such a draw a *seed*. Every insurer in the fourth seat runs on the same 200 seeds, so two insurers in the fourth seat are compared on the same draws. The market at the values of Table A.1 is *the main market*, which Section 6 varies. Growth is measured on the surplus at the close of the measured years, the *terminal surplus*. The surplus the insurer held when the measured years began is its *endowment*. Terminal surplus is compared with the endowment grown at the *return on funds*, the yield that surplus and reserves earn, so that a surplus that earns that return and nothing else has not grown. Growth is stated as a rate of growth a year or as a multiple of the endowment so grown. A difference between two insurers in the fourth seat is stated in percent of terminal surplus: by how many percent the one ends with more surplus than the other. A difference is *resolved*, statistically distinguishable from zero, where its t statistic, the mean over its standard error across seeds, exceeds 2 in magnitude. The market is stylized, and the Kelly decision is the paper's construction, not an observed insurer's practice; the inputs follow industry practice where practice is documented, and each is listed with its grounding (Appendix A).

The paper has three findings. First, in the main market learning the loss adds more growth than anything else the paper measures. The benchmark grows at 1.93 percent a year, a multiple of 1.71 over the measured years. The plan route on the learned estimate grows at 3.39 percent a year, a multiple of 2.55: learning the loss adds 49.2 percent to terminal surplus. Learning an account's claim frequency is worth more than learning the size of its claims (Section 4).

Second, in the main market, on the plan route, learning the loss pays on new business, and on renewals the rating plan's estimate grows surplus more. On a renewal the insurer holds the account's claim file, which both the rating plan, through the experience modification, and the learned estimate read (§4.3). The plan route with learning on new business adds 62.4 percent of terminal surplus over the benchmark, and learning on renewals alone takes away 15.1 percent. The plan route with learning on new business grows at 3.71 percent a year, a multiple of 2.77, the highest growth of any insurer the paper reports in the fourth seat in the main market, 8.8 percent of terminal surplus above the plan route on the learned estimate (Section 4).

Third, in the main market, what the Kelly decision adds depends on the estimate it prices. On the rating plan's estimate the Kelly route ends with 21.8 percent more terminal surplus than the plan route. On the learned estimate the Kelly route and the plan route end with equivalent growth, at a margin set after the difference was measured: the 90 percent interval of their difference, 0.2 to 2.0 percent of terminal surplus, lies inside an *equivalence margin* of 2.8 percent of terminal surplus either side of zero, 0.10 percentage points of growth a year, and the paper counts two growths whose interval lies inside the margin as equivalent (§5.1). With learning on new business the Kelly route trails the plan route by 9.6 percent, and it trails from the first measured year (Section 5).

The plan route with learning on new business was also run at other values of four dimensions of the market, one at a time: the Bureau insurer's refiling threshold; the noise of an underwriter's read of an account (the read is the judgment that sets the schedule modification); the *filed band*, the range either way within which schedule rating may move a price, 25 percent in the main market and 40 percent at the wider band; and the buyer's responsiveness to price. It was also run in four further markets (Section 6). The Kelly route, with learning on new business and on the learned estimate, was run at the wider band, at the two other values of the buyer's responsiveness, and in the four further markets. The plan route with learning on new business grows more than the benchmark wherever the two were run (under the severity trend, not at the level that allows for the number of tests, §7.2). It grows more than the plan route on the learned estimate at every other value run of the main market's dimensions (at the lower read noise, not at that level) and with smaller insurers. In a market with a second insurer that learns, and in one with a severity trend, a yearly growth in claim sizes that no insurer projects, the plan route on the learned estimate grows more (under the trend, not at that level); in the market of one tenth of the accounts the difference is not resolved. With learning on new business the Kelly route trails the plan route wherever the difference is resolved (under the severity trend, not at that level). On the learned estimate the Kelly route's difference from the plan route changes sign from setting to setting: it runs from −19.7 percent of terminal surplus at the wider band to +3.1 percent with the less responsive buyer. In the four further markets the Kelly route on the learned estimate leads (with smaller insurers, not at that level), by up to 12.3 percent of terminal surplus with one tenth of the accounts, or, beside a second insurer that learns, ends within the equivalence margin, which was set after the difference there too was measured.

The limits are stated in full in §7.2. The market is one line of business, with no loss trend except in the one market run with one, no underwriting cycle, and no reinsurance. Its conventional insurers respond only through their own rules to an insurer that learns. The full grid is run in the main market alone. At the main market's surplus the criterion's weight on risk, the decision's variance term, is equivalent to a charge of 0.0641 percent of the expected loss of an account at the bureau's scale of claim sizes, and removing it is shown not to move the growth of the Kelly route on the learned estimate by more than 0.010 in the logarithm. In the market of one tenth of the accounts the term is 10.0 times its size in the main market, and removing it lowers that growth by 2.3 percent of terminal surplus (§D.4; not at that level). The paper does not run a market whose

accounts are large against an insurer's surplus. The horizon is known to the insurer, and growth is read at it. And every difference is measured within one model over 200 seeds, so its t statistic measures the precision of the simulation, not of an observed market.

The paper draws on three literatures. The first is the Kelly criterion of Kelly and Breiman. Its second-order form, the approximation that keeps the mean and the variance of an outcome, weighs an account's expected margin against the variance of its loss, and the criterion, not the insurer, sets how much the variance counts.[3] The second is the theory of insurance pricing under asymmetric learning. In it the insurer that holds an account learns about it what competitors do not. Where the insurer cannot bind itself to its later prices, it prices the first year low and recovers the shortfall through the margins of later renewals on the accounts that stay; that literature calls the low first price "lowballing."[4] In the market of this paper the holder reads its own claim file where a competitor reads a broker's loss run, and every contract is annual; the paper does not model the equilibrium that literature derives. In the main market the Kelly route on the learned estimate quotes new business below and renewals above its undeviated price (§5.2). The empirical counterpart is the measured difference between new and renewal business: in each of the 25 books of business that Wu and Lin examine the loss ratio of new business is above that of renewals.[5] The third is the study of pricing algorithms in competition, which asks whether algorithms that learn, or that respond to rivals' prices, raise prices.[6] That question is about the level of prices among sellers whose algorithms meet one another; the question here is what an objective and the learning are worth to one insurer among insurers that do not learn as it does.

Section 2 states the approach: the objective, the four things the decision needs and how each enters the price, and what learning supplies. Section 3 describes the market. Section 4 reports what learning the loss is worth, what is learned, and where it is used: in the main market the claim frequency is worth more than the claim size, and, on the plan route, learning the loss pays on new business and not on renewals. Section 5 reports what the Kelly decision adds on each estimate: in the main market the Kelly route grows more than the plan route on the rating plan's estimate, ends within the equivalence margin of the plan route on the learned estimate, a margin set after the difference was measured, and falls short of the plan route with learning on new business. Section 6 reports the other settings and the further markets in which learning on new business and the two routes were run. Section 7 states what follows for an insurer, the limits, and the conclusion. The appendices carry the market's assumptions and its construction in full, the rating plans, the capital

[3] Kelly, "Interpretation of Information Rate"; Breiman, "Optimal Gambling Systems"; Edward O. Thorp, "The Kelly Criterion in Blackjack Sports Betting, and the Stock Market," in *Handbook of Asset and Liability Management*, vol. 1, ed. Stavros A. Zenios and William T. Ziemba (Amsterdam: North-Holland, 2006), 385–428.

[4] Howard Kunreuther and Mark Pauly, "Market Equilibrium with Private Knowledge: An Insurance Example," *Journal of Public Economics* 26, no. 3 (1985): 269–288; Tore Nilssen, "Consumer Lock-in with Asymmetric Information," *International Journal of Industrial Organization* 18, no. 4 (2000): 641–666; Thomas de Garidel-Thoron, "Welfare-Improving Asymmetric Information in Dynamic Insurance Markets," *Journal of Political Economy* 113, no. 1 (2005): 121–150; Ruo Jia and Zenan Wu, "Insurer Commitment and Dynamic Pricing Pattern," *Geneva Risk and Insurance Review* 44, no. 1 (2019): 87–135; and, for evidence consistent with the theory, Stephen P. D'Arcy and Neil A. Doherty, "Adverse Selection, Private Information, and Lowballing in Insurance Markets," *Journal of Business* 63, no. 2 (1990): 145–164.

[5] Cheng-Sheng Peter Wu and Hua Lin, "Large Scale Analysis of Persistency and Renewal Discounts for Property and Casualty Insurance," *Casualty Actuarial Society E-Forum* (Winter 2009): 398; Stephen P. D'Arcy and Neil A. Doherty, "The Aging Phenomenon and Insurance Prices," *Proceedings of the Casualty Actuarial Society* 76 (1989): 24–44.

[6] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello, "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review* 110, no. 10 (2020): 3267–3297; Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy* 132, no. 3 (2024): 723–771; Zach Y. Brown and Alexander MacKay, "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics* 15, no. 2 (2023): 109–156.

standard, the Kelly decision and the estimators with their equations, the Kelly route's computation step by step, the registry of predictions and outcomes, the quote data, and the further markets.

2 Growth as the Objective, and Learning as Its Source

This section states the objective, the price it gives for one account, and what an insurer learns from its own records to set that price.

2.1 The Objective, and the Decision It Gives

The Kelly criterion takes the long-run growth of surplus as the objective: the insurer chooses its prices so as to maximize the expected logarithm of its surplus at the horizon, which is the same as maximizing the expected mean geometric rate at which the surplus grows from year to year. The criterion's property holds in repeated independent bets on known odds (Section 1); an insurer does not know the odds, and the paper measures what the objective is worth to an insurer that has to estimate them. Appendix D states the approach of this section with its equations and the estimators in full.

For one account the objective reduces to a rule. Writing an account at a premium changes the surplus by the premium net of expenses less the account's loss. For an account small against the surplus, the expected change in the logarithm of surplus is approximately (§D.1.2) the expected margin as a share of surplus, less half the variance of the loss divided by the square of the surplus: a mean–variance rule, in which the objective, and not the insurer, sets how much the variance counts. The reserve held against the expected loss earns the return on funds until the claims are paid. And in a market where the account can be kept, an account written today is likely to be held in the years to come, so the margins of those years count. The paper discounts them in two ways (Table 1). The *growth discount* discounts them at the insurer's own rate of growth of surplus, the return on funds included, the rate the objective implies. The *funds discount* discounts them at the return on funds. The expected increment to the logarithm of surplus in the paper's construction is

$$G(P) \;=\; \frac{n - \hat{x}\,(1 - c_x) + \bar{n}\,V}{S} - \frac{\hat{v}}{2S^2}, \tag{1}$$

where P is the premium quoted and n that premium net of expenses; $\bar{n}$ is the premium net of expenses at the undeviated price, from which the decision deviates; $\hat{x}$ and $\hat{v}$ are the account's estimated expected loss and the variance of its loss; c_x is the return on funds the reserve earns before the loss is paid, as a share of the expected loss, from which it is netted; V is the value of keeping the account per unit of $\bar{n}$; and S is the surplus. The Kelly decision is the premium that maximizes the chance of winning the account at that premium times the increment it would earn. In words, the expected increment is the premium net of expenses, less the expected loss net of what the reserve held against it earns before the loss is paid, plus the value of keeping the account, all as a share of surplus, less half the variance of the loss over the square of surplus. A higher premium adds margin and wins the account less often, and the decision is the premium at which the two are best traded. It gives both answers an insurer needs at a quote: the price and, by the floor below, whether the insurer wants the account at that price. It needs four things of an account. The expected loss $\hat{x}$ is what a rating plan also estimates. The variance $\hat{v}$ enters through the criterion's weight on risk, the variance term of Equation 1: of two accounts with one expected loss, the one whose loss is more variable is charged at least as much, and an insurer with less surplus charges at least as much for

the same variability. The value of keeping the account per unit of $\bar{n}$, V, is the margin per unit of net premium of each later renewal times the chance of reaching it, discounted, summed over the years to the horizon; the paper calls it the account's *lifetime value*. It counts the renewals after the policy quoted, not the policy quoted, and only to the horizon. It is a value to the insurer, counted in its objective, and not a discount granted to the buyer, on renewal or on new business. And the chance of winning at each premium, a renewal kept or, on new business, an account won among those that go to market, is the insurer's own *hold curve*, fitted to the quotes it won and lost.

A markup today changes today's margin and the chance of winning the account. To a first approximation it does not change next year's margin: next year's quote comes from next year's rule applied to next year's estimate, and the buyer who shops compares the holder's renewal quote with competitors' quotes, not with this year's markup. This is argued from the market's rules; the paper does not measure it (§D.1.5). The chance of winning multiplies the whole increment, so what a markup does to that chance is already counted there; the lifetime value itself therefore carries no markup and enters at the undeviated price, $\bar{n}$ in Equation 1: the *anchored lifetime value*. The other form, the *quote-anchored lifetime value*, lets the lifetime value rise with the premium quoted, n in place of $\bar{n}$, and so counts, for a markup, a value that by that argument the markup does not buy; the paper measures what each form is worth (§5.1, Appendix F). The undeviated price is below the Kelly route's own renewal prices on average, so the anchored lifetime value understates the lifetime value (§7.2). The decision also sets a floor, the premium at which the increment with the quote-anchored lifetime value is zero, below which the insurer does not write; in the main market it binds no quote of the Kelly route on the learned estimate (Table G.1).

2.2 Learning from What an Insurer Holds

A rating plan prices an account's history through one number, the experience modification on its capped losses. The learning of this paper has two parts, each read from data the insurer holds: learning the loss, each account's loss distribution estimated from the whole file, and learning the hold curve, the chance of winning fitted from the broker's figures, by how much its lost quotes lost. First, the insurer tracks two traits of each account: its claim frequency relative to its class, and the scale of its claim sizes. It reads them from the file it has: its own claim file on an account it holds, and the broker's loss run on an account it is asked to quote. It weights the years of the file by credibility, each year by how much its claims carry and by how far back it lies. And it estimates from the claim files and loss runs it has read, rather than assumes, how long an account's departure from its class persists. The expected loss and the variance follow from the tracked frequency and scale, and the variance is the quantity a plan does not hold.

Second, the insurer fits its own hold curve: the chance of winning at each deviation from the undeviated price. It fits the curve from the only information it has on competitors' prices: the broker's figure, by how much a quote lost where the account went to market, and which of its quotes held their accounts. The fit allows for the quotes that bound, whose competing prices the insurer never sees, and the decision averages the chance of winning over the uncertainty of the fit. Beside the two parts, the lifetime value is read from the insurer's own book over a window of years: the retention and the margin at a first renewal and at later renewals, and the insurer's own growth of surplus, the return on funds included. The four things of §2.1 come from these parts and that book: the expected loss and its variance from the first, the chance of winning from the second, and the value of keeping the account from the insurer's own book.

On the plan route the learned expected loss takes the place of the class loss cost with its experience modification, and the underwriter's schedule modification multiplies it as it does the rating plan's estimate. The plan's own rules then turn it into a quote. The rate level is moved by the loss-ratio method, with competitors' winning prices as the complement of credibility, the figure that receives the weight not given to the insurer's own experience. The *schedule credit* is the credit each insurer gives new business from the winning prices it has seen (Section 3). On the Kelly route the Kelly decision chooses a deviation from the undeviated price on the hold curve. The deviation, the schedule modification, and the schedule credit together stay inside the filed band, 25 percent either way in the main market and 40 percent at the wider band of Section 6, and a zero deviation on every account reproduces the plan route. The plan route uses the first part of the learning alone, the learned expected loss; the Kelly route uses both, with the variance entering through the decision. Each prices at its own rate level, moved each year by the loss-ratio method on its own book (Appendix E states the Kelly route's computation step by step).

Learning on new business uses the learned expected loss on new business alone. On a renewal the insurer prices on the rating plan's estimate, the class loss cost with the experience modification computed from its own claim file. The rest of each route is unchanged: on the plan route the plan's rules turn the estimate into a quote, and on the Kelly route the decision deviates from the undeviated price, with the learned variance on both kinds of business.

The Kelly route on the rating plan's estimate prices the plan's expected loss. The variance it prices is that of a loss at the bureau's claim frequency for the account's class and size times the account's experience modification, with the class's claim sizes on the bureau's curve of claim sizes (§D.2). As on the other two estimates, it fits its own hold curve to the quotes it won and lost and reads the lifetime value from its own book.

3 A Stylized Market

The learning and the decision are measured in a market of 4,000 accounts of middle-market general liability, written by four insurers without reinsurance; Figure 1 maps one year of it, and Appendix A states every assumption with its value, its grounding, and the variations of the market the paper runs (Section 6; Appendix F). Each account has a size and two traits, one of claim frequency and one of claim size, each with a permanent part and a part that drifts from year to year; claims occur, are reported with a lag, and are paid at settlement, with an open claim carried at a case reserve with error. The claim sizes follow the distribution increased-limits ratemaking uses, a mixed exponential, scaled so that the book's mean claim at the occurrence limit is the industry's published general liability severity. Limits apply per occurrence and in the aggregate, and loss adjustment expense is in proportion to the indemnity paid (Appendix A).

Three sources of information exist, and they are not equally shared. A bureau publishes, for each of five classes, the market's pooled loss costs and claim frequency of past policy years. An insurer asked to quote receives the broker's loss run, the account's claims of the past five policy years whoever wrote them, and its underwriter forms a read of the account of its own, which sets the schedule modification (Appendix B). The insurer that holds an account reads its own claim file for every year it has held it, current to the quarter and with a smaller reserve error than the broker's run carries. The holder therefore knows more about its accounts than a competitor invited to quote does, and the longer it has held an account the more; nothing else in the market favors the holder except the buyer's costs of shopping and switching.

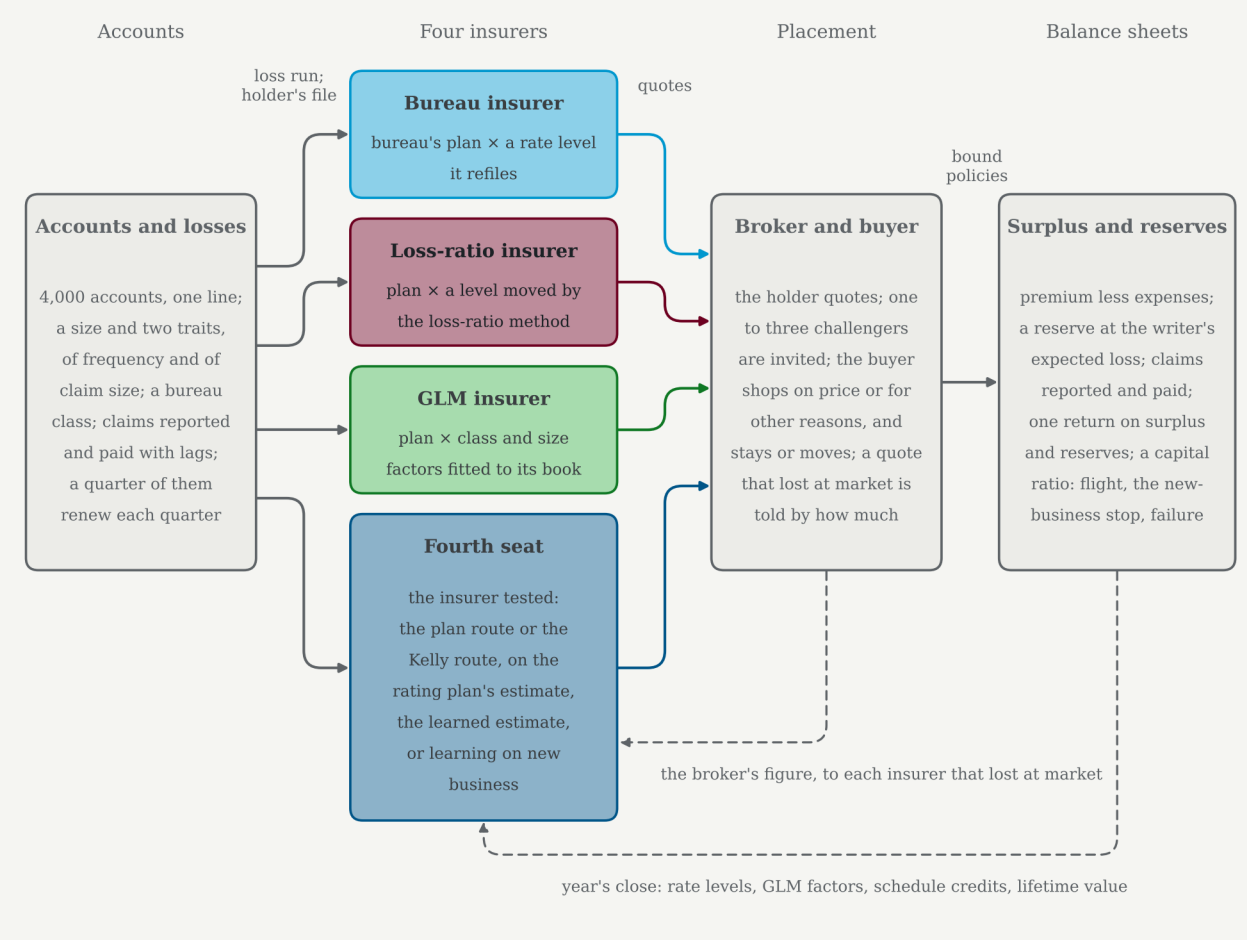


Figure 1: One year of the market. Each quarter the holder quotes, and so do the challengers, the competitors the broker invites; the buyer chooses; an insurer that lost at market is told the broker's figure; and the policy is booked. At the year's close the rate levels, the GLM insurer's factors, the schedule credits, and the lifetime value are re-estimated, and capital is tested, with broker flight: brokers withhold new-business invitations from an insurer at a low capital ratio (§A.5). Solid arrows are flows within a quarter, dashed arrows what returns to an insurer. The fourth seat holds each insurer of Table 1 in turn. Source: authors' construction.

Three conventional insurers compete in every market the paper runs but one, in which a second insurer that learns takes the GLM insurer's seat (Section 6). Each prices by the rating plan, the bureau's class loss cost times the experience and schedule modifications times its own rate level, with a provision for expenses near the industry's published expense ratio and for a stated profit (Appendix A). Each sets its level in its own way. The Bureau insurer refiles its rate level when the indication on its own experience departs from the plan by a stated threshold. The Loss-ratio insurer moves its level every year by the loss-ratio method, by how its book's losses ran against what it was priced to earn, with the level of competitors' winning prices as the complement of credibility. The GLM insurer moves its level by its book's losses against the losses its own relativities imply, credibility-weighted, and adds relativities by class and size band from a generalized linear model fitted to that book. No insurer files a load on new business. Each gives a schedule credit on new business in the larger size bands, set from the winning prices it has seen and inside the filed band, 25 percent either way in the main market. In the two books of general liability that

Wu and Lin examine, new policies receive schedule debits about half as often as renewals and schedule credits as often: new business is priced below renewals through the plan's own flexible components.[7] Beyond these rules, the three do not respond to a competitor that learns. The main market's filed band is taken from a Florida rule on schedule rating, which excludes competitors' prices from what a schedule may rate (§7.2). Appendices A and B state the plans and the rules in full.

An account renews once a year; its holder quotes, and the broker invites one or more *challengers*, competitors asked to quote, more for a larger account. The chance that the buyer goes to market rises with the gap between the holder's renewal quote and the lowest competing quote, against a cost of shopping, and the chance that it then switches rises with the same gap against a cost of switching; with probability 0.10 an account is remarketed for reasons unrelated to price. The two costs were calibrated, on other seeds, so that 14 percent of the accounts change insurer each year in the market with the benchmark in the fourth seat, a market of four conventional insurers. The target is half the share of middle-market companies that a 2007 survey of the segment found changing a property–casualty or employee-benefits insurer over two years; halving it counts no company changing twice. The calibration gives a premium retention between two retentions of commercial business published in 2026 (Appendix A). After placement each insurer that quoted on an account that went to market and lost is told the difference between its quote and the bound premium, rounded, *the broker's figure*; the insurer that bound is told nothing. Every insurer starts the measured years with the same endowment, set so that its opening book stands at the industry's aggregate capital ratio under the risk-based-capital standard. Its ratio of written premium to surplus is then 0.338, against the industry's 0.773, because the standard charges reserves as well as premium and in steady state a book of this one line holds 2.69 dollars of reserves for each dollar of premium, against the industry's 1.06. Surplus and reserves earn the return on funds, the industry's published investment yield. Capital matters at three thresholds (Appendix C): a broker withholds new-business invitations from an insurer whose capital ratio falls into the band at which regulators watch it, *broker flight*; an insurer whose written premium exceeds three times its surplus stops quoting new business; and an insurer below the level at which regulators take control fails and is replaced. No insurer fails in any seed of any market the paper runs.

After a build-up of the books and files (§A.5), the market runs 10 practice years before the 28 measured years begin, so that every insurer opens the measured years with estimators that have run and a book shaped by its own pricing. An insurer's *growth* is the logarithm of its terminal surplus over its endowment, net of the return on funds, so that a surplus that earned that return and nothing else has a growth of zero; the paper states it as a rate of growth a year and as the multiple of the endowment grown at the return on funds. Tables 1 and 3 give growth as the logarithm. The benchmark's growth of 0.534 in the logarithm is a multiple of 1.71, or 1.93 percent a year over the 28 measured years. The exponential of a difference in the logarithm, less one, is the percent by which one terminal surplus exceeds the other: the lead of the plan route on the learned estimate over the benchmark, 0.400 in the logarithm, is 49.2 percent. A difference between two insurers in the fourth seat is taken within each seed and averaged, and its t statistic states how precisely the simulation resolves it. The paper states a difference in growth as a percentage of terminal surplus, computed from the unrounded means; the tables carry each difference's t statistic, and a difference the text states without qualification is resolved. Every difference is a total effect: what changes in one seat moves the other insurers' books, rate levels, and capital, and the difference carries their response.

[7] Wu and Lin, "Large Scale Analysis," 399, 404–405.

4 What Learning the Loss Is Worth

Learning the loss is the largest effect in the main market. This section measures it at the plan's pricing: the insurer in the fourth seat prices its estimate through the rating plan, as the Loss-ratio insurer does, and only the estimate changes. Table 1 gives 14 insurers in the fourth seat, and Table 2 further insurers that take learning the loss apart, all on the same 200 seeds with the three conventional insurers the same in each.

4.1 The Learned Estimate

The benchmark grows at 1.93 percent a year, a multiple of 1.71 over the measured years. It and the Loss-ratio insurer beside it, one rule in two seats, end 2.9 percent of terminal surplus apart (t = +1.60), a difference that is not resolved. The plan route, the learned estimate priced as the Loss-ratio insurer prices, grows at 3.39 percent a year, a multiple of 2.55. Learning the loss adds 49.2 percent to terminal surplus. The three conventional insurers grow less beside the plan route than beside the benchmark (Figure 2). In the market with the plan route in the fourth seat, the plan route ends with 71.8, 84.4, and 68.0 percent more terminal surplus than the Bureau insurer, the Loss-ratio insurer, and the GLM insurer beside it (Appendix F, Table F.1).

Table 1: The insurers in the fourth seat, means over 200 seeds, by the estimate each prices on and the way it prices. Growth is the logarithm of the fourth seat's terminal surplus over its endowment, net of the return on funds (Section 3); the rate a year, in percent, and the multiple of the endowment grown at the return on funds restate the mean logarithm. The share is the fourth seat's share of the accounts at the close, and the final column the share of the market's accounts changing insurer a year. On each estimate the Kelly route's first row is the main construction: the anchored lifetime value, the funds discount (§2.1), and the pessimism variant, which lowers the chance of winning the insurer assumes at every deviation (§D.3). On the learned estimate and with learning on new business its variations follow it, each named by its choices: with or without the pessimism variant, the funds discount or the growth discount (at the insurer's own growth of surplus), and the quote-anchored lifetime value (§2.1). Source: authors' calculations.

The fourth seat	Growth, logarithm	Rate a year, %	Growth multiple	Share of accounts	Changing insurer a year
The rating plan's estimate					
The plan route: the benchmark	+0.534	1.93	1.71	0.223	0.141
The Kelly route	+0.732	2.65	2.08	0.231	0.241
The learned estimate					
The plan route	+0.935	3.39	2.55	0.672	0.163
The Kelly route	+0.945	3.43	2.57	0.494	0.263
Without pessimism, funds discount	+0.945	3.43	2.57	0.488	0.263
With pessimism, growth discount	+0.941	3.42	2.56	0.461	0.270
Without pessimism, growth discount	+0.939	3.41	2.56	0.452	0.269
Quote-anchored, without pessimism, funds discount	+0.836	3.03	2.31	0.308	0.258
Learning on new business					
The plan route	+1.019	3.71	2.77	0.655	0.198
The Kelly route	+0.918	3.33	2.50	0.296	0.270
Without pessimism, funds discount	+0.911	3.31	2.49	0.284	0.271
With pessimism, growth discount	+0.919	3.34	2.51	0.290	0.280
Without pessimism, growth discount	+0.908	3.30	2.48	0.281	0.280
Quote-anchored, without pessimism, funds discount	+0.796	2.88	2.22	0.193	0.290

The new business the benchmark wins has a true expected loss of 1.012 times its own estimate, consistent with the "winner's curse": an insurer wins more often where its estimate is low.[8] The new business the plan route wins has a true expected loss of 0.927 times its estimate; the paper does not measure why the ratio falls below one. The insurer is paid on its estimate, and with the learned estimate it does not win mainly where its estimate is low. The plan route also holds 67 percent of the market's accounts at the close, a book at which the capital standard begins to bind: brokers withhold 964 new-business invitations over the 28 measured years and 200 seeds, where in the market with the benchmark they withhold none. The paper does not separate what the size of a book contributes to growth from what its margin contributes (§7.2). Figure 2 gives every insurer's growth by seed in the two markets.

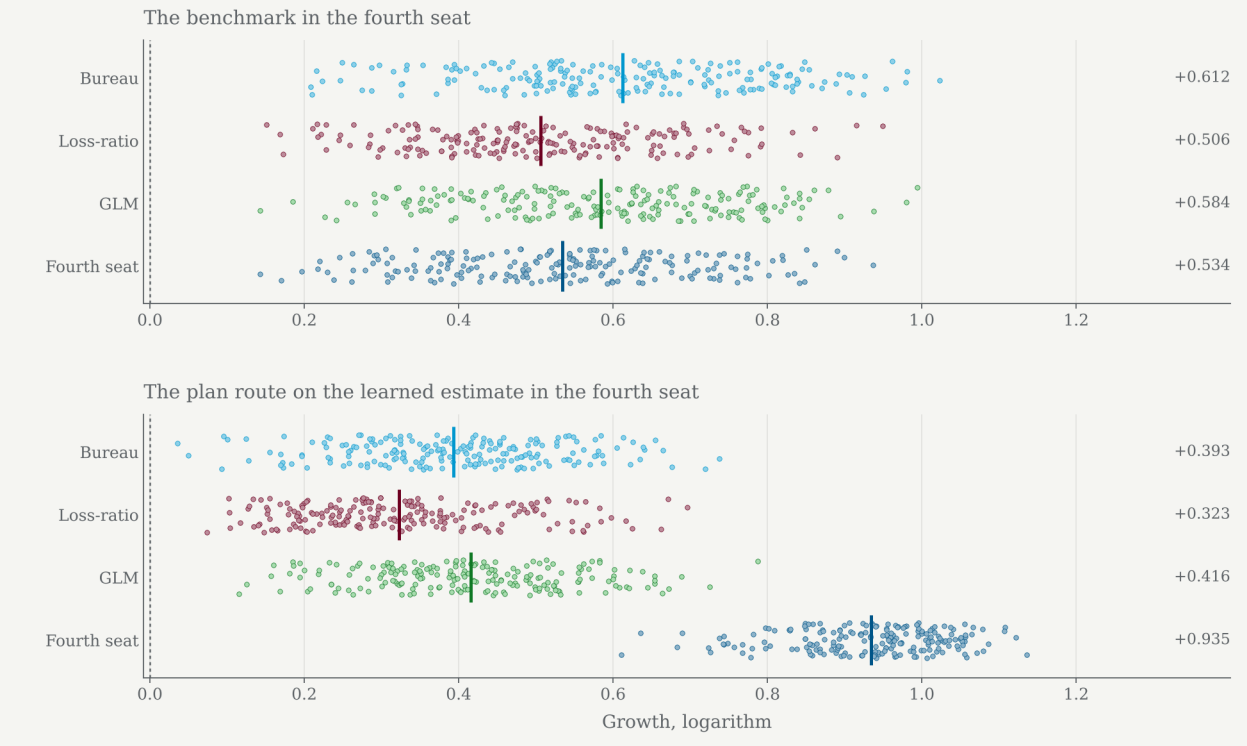


Figure 2: The three conventional insurers grow less beside the plan route on the learned estimate than beside the benchmark. Growth by insurer and seed with the benchmark in the fourth seat (upper panel) and with the plan route on the learned estimate there (lower panel), over 200 seeds. Each insurer's row of points gives its growth, the logarithm of its terminal surplus over its endowment grown at the return on funds (Section 3), on one horizontal scale, a vertical bar marking the mean, stated at right. In all figures the Bureau insurer is bright blue, the Loss-ratio insurer maroon, the GLM insurer green, and the fourth seat deep blue. Source: authors' calculations.

Learning the loss adds to growth at four other settings of two dimensions of the market: the Bureau insurer's refiling threshold, 0.10 in the main market and 0.05 and 0.15 at the other settings; and the noise of an underwriter's read of an account, the read noise, as a multiple of the spread of

[8] Edward C. Capen, Robert V. Clapp, and William M. Campbell, "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology* 23, no. 6 (1971): 641–653; Paul R. Milgrom and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica* 50, no. 5 (1982): 1089–1122.

the account's traits, 0.7 there and 0.35 and 1.05 at the other settings. The gain is resolved at each setting (Appendix F, Table F.1).

4.2 What Is Learned

The learned estimate tracks two traits of each account: its claim frequency relative to its class and the scale of its claim sizes. To see what each is worth, the insurer in the fourth seat learned one trait and held the other at its class value, the value the estimator gives an account with an empty file. Table 2 gives the contrasts. Its first block starts from the class value of both traits, with no experience modification, a lower base than the benchmark of §4.1, which uses the experience modification. Against the class value of both traits, learning the frequency alone adds 41.7 percent to terminal surplus, and learning the claim size alone adds 5.8 percent. Learning both adds 111.9 percent, more than the two alone: each trait is worth more when the other is learned too. The claim frequency is worth more than the claim size, alone and with the other trait learned.

Table 2: What learning the loss is worth, by what is learned, by how an account is read, and by where learning is used; and learning on new business against the learned estimate: the fourth seat's growth with one insurer there less another, paired by seed over 200 seeds, at the plan's pricing. Each row gives the mean in the logarithm, its t statistic, the seeds above zero, and the difference in percent of terminal surplus. The class value of a trait is the estimator's value for an account with an empty file. Source: authors' calculations.

The contrast	Mean, logarithm	t	Seeds above zero	Percent of terminal surplus
What is learned, against the class value of both traits				
The claim frequency alone	+0.349	+45.71	199	+41.7
The claim size alone	+0.057	+14.05	174	+5.8
Both traits	+0.751	+100.28	200	+111.9
The frequency, against the claim size	+0.292	+36.25	200	+33.9
How an account is read				
The experience modification, against the class loss cost	+0.345	+29.79	198	+41.2
The moments of the traits given, against the moments estimated	+0.035	+5.52	120	+3.5
Where learning is used, against the benchmark				
New business alone	+0.485	+38.92	199	+62.4
Renewals alone	−0.164	−12.85	34	−15.1
Both	+0.400	+31.10	197	+49.2
Learning on new business, against the learned estimate				
New business alone, against both kinds	+0.084	+12.48	166	+8.8

The rating plan prices an account's history through the experience modification. Against the class loss cost alone, the experience modification adds 41.2 percent of terminal surplus. The insurer also estimates, from the claim files and loss runs it has read, the second moments of the traits, among them how far an account's traits stay away from its class over time. Given in their place the moments that a large sample of the loss process yields, the insurer ends with 3.5 percent more terminal surplus.

4.3 Where Learning Is Used

An insurer quotes two kinds of business. On new business it quotes an account it does not hold, on the broker's loss run. On a renewal it quotes its own account, on its own claim file. The insurer in the fourth seat used learning on one kind alone and the rating plan's estimate on the other. Learning on new business alone adds 62.4 percent of terminal surplus over the benchmark. Learning on renewals alone takes away 15.1 percent. On renewals the rating plan's estimate grows surplus more than the learned estimate.

The selection of the business each insurer keeps shows the same. The plan route on the learned estimate keeps renewals that carry a true expected loss of 1.027 times its estimate. For the plan route with learning on new business, the renewals it keeps, priced on the rating plan's estimate, carry 0.997 times its estimate, and its new business 0.989 times. An insurer is paid on its estimate, so an estimate that runs low on the accounts it keeps costs it margin. The paper does not measure why learning on renewals falls short of the rating plan's estimate beyond the selection of the renewals kept.

The plan route with learning on new business grows at 3.71 percent a year, a multiple of 2.77: the highest growth of any insurer the paper reports in the fourth seat in the main market (Tables 1 and 2). It adds 8.8 percent of terminal surplus to the plan route on the learned estimate, with 166 of 200 seeds above zero (t = +12.48). It holds 65 percent of the market's accounts at the close, and 20 percent of the market's accounts change insurer a year.

5 The Kelly Decision on Each Estimate

The Kelly decision prices an estimate as a deviation from the undeviated price of that estimate, inside the filed band (Section 2). It was run on each of the three estimates. The main construction makes three choices: the anchored lifetime value (§2.1); the margins of later renewals discounted at the return on funds, the funds discount; and the pessimism variant, in which the insurer lowers the fitted center of its hold curve, the deviation at which the fitted chance of winning is one half, by one posterior standard deviation, so that the chance of winning it assumes is lower at every deviation than the fit's mean gives (§D.3). On the learned estimate the variant's effect on the Kelly route's growth is not resolved (0.0 percent of terminal surplus, t = +0.03; Table F.2). The Kelly route on the rating plan's estimate makes the same three. Table 1 gives the variations on the learned estimate and with learning on new business. Table 3 sets the six insurers in the fourth seat in a grid, the estimate by the pricing, with the Kelly route less the plan route on each estimate. Figure 3 draws that difference for learning on new business.

Table 3: The grid: the fourth seat's growth by the estimate it prices on and the way it prices, means over 200 seeds; and the Kelly route less the plan route on the same estimate, paired by seed, in percent of terminal surplus, with its t statistic and its reading at the equivalence margin of 2.8 percent of terminal surplus, a margin set after the difference on the learned estimate was measured. *Within* means that the 90 percent interval of the difference lies inside the margin. On each estimate the Kelly route is the main construction. Source: authors' calculations.

	Growth, logarithm		The Kelly route less the plan route		
The estimate	The plan route	The Kelly route	Percent	t	At the margin
The rating plan's estimate	$+0.534$	$+0.732$	$+21.8$	$+18.99$	outside
The learned estimate	$+0.935$	$+0.945$	$+1.1$	$+1.97$	within
Learning on new business	$+1.019$	$+0.918$	-9.6	-14.49	outside

5.1 What the Decision Adds

On the rating plan's estimate the Kelly route ends with 21.8 percent more terminal surplus than the plan route (t = +18.99). On that estimate the Kelly route prices the plan's expected loss with its own hold curve and lifetime value (§2.2). The Kelly route quotes renewals above its undeviated price and new business below it (Appendix F, the eighth registration).

On the learned estimate the two routes end the measured years with equivalent growth. The Kelly route grows at 3.43 percent a year, a multiple of 2.57, and the plan route at 3.39 percent, a multiple of 2.55. The Kelly route less the plan route is 1.1 percent of terminal surplus (t = +1.97). Read at the close of each measured year, the Kelly route trails the plan route by a resolved difference in every year from year 1 through year 23, deepest in year 10, where the difference is -5.0 percent of that year's surplus. The difference then narrows: from year 22 through the close its 90 percent interval lies inside the margin, and the difference ends above zero; the paper does not divide that narrowing between the decision and the approach of the horizon (§7.2). A difference that is not resolved is not a finding of equality, so the paper tests equivalence against a margin: 0.028 in the logarithm over the 28 measured years, 2.8 percent of terminal surplus or 0.10 percentage points of the rate of growth a year. Two growths are *equivalent* where the 90 percent interval of their difference lies inside the margin either side of zero. The margin was set with this difference already measured: a tenth of a percentage point of growth a year, below the smallest effects on which the registered findings made before it rest (Appendix F). The interval runs from 0.2 to 2.0 percent of terminal surplus. It lies above zero and inside the margin: a 90 percent interval excludes zero at $|t| > 1.645$, while a difference is resolved only at $|t| > 2$.

With learning on new business the Kelly route trails the plan route by 9.6 percent of terminal surplus (t = -14.49), beyond the margin. It grows at 3.33 percent a year, a multiple of 2.50. Learning on new business in place of the learned estimate adds 8.8 percent to the plan route's terminal surplus (t = +12.48) and takes 2.7 percent from the Kelly route's (t = -3.52, not at the level that allows for the number of tests, §7.2). In the main market, the better the estimate, measured by the growth it gives at the plan's pricing, the less the Kelly decision adds: the Kelly route ends with 21.8 percent more terminal surplus than the plan route on the rating plan's estimate, within the equivalence margin of the plan route on the learned estimate, and 9.6 percent short of the plan route with learning on new business. The shortfall is not the main construction's alone: each of its four variations trails the plan route with learning on new business, by 9.5 to 20.0 percent of terminal surplus (Table 1 gives their growth; Appendix F).

In the main construction the lifetime value is anchored at the undeviated price, the anchored lifetime value (§2.1). The variation with the quote-anchored lifetime value, without pessimism and on the funds discount (Table 1), lets the lifetime value rise with the price quoted; on the learned estimate it ends 9.4 percent of terminal surplus below the plan route. The same variation with the anchored lifetime value ends 11.6 percent of terminal surplus above the quote-anchored one; that is the worth of the anchoring (Appendix F).

5.2 Where the Two Routes Differ

The two routes quote differently, and with learning on new business they build different books (Tables 1 and 4). On the learned estimate the Kelly route quotes new business 8.4 percent below its undeviated price and renewals 12.6 percent above it. That is a low first price and a higher renewal price, the shape of the lowballing that the theory of insurance pricing under asymmetric learning predicts of an insurer that knows its own accounts better than its competitors do.[9] With learning on new business the Kelly route quotes new business 1.2 percent above its undeviated price and renewals 4.3 percent above it on average, with 14 percent of its renewal quotes within one grid step, 0.25 percent, of the band's upper end. The undeviated price is the Kelly route's own, at its own rate level and, on new business, with its own schedule credit; it is not the price the plan route quotes. Each mean deviation is averaged over every quote, won or lost; the premium over the true expected loss in Table 4 is over the business written. The Kelly route holds 30 percent of the market's accounts at the close against the plan route's 65 percent, and 27 percent of the market's accounts change insurer a year beside it, against 20 percent beside the plan route.

Read at the close of each measured year, the Kelly route trails the plan route by a resolved difference in every year from year 1 through year 28, deepest in year 17, where the difference is -10.4 percent of that year's surplus (Figure 3). At the close it trails in 173 of 200 seeds. On average along a path its growth is the steadier of the two, the standard deviation of its yearly increments of growth the smaller ($t = -7.45$ for the difference in volatility), and it is the less steady in 56 of the 200 seeds. Across seeds its terminal surplus is the more dispersed, the standard deviation of its growth at the close 1.272 times the plan route's.[10] The paper does not establish the cause of the shortfall; Table 4 gives the differences in price and in book that come with it.

[9] D'Arcy and Doherty, "Adverse Selection"; Jia and Wu, "Insurer Commitment," 89.

[10] The t of two paired standard deviations, Table 4: E. J. G. Pitman, "A Note on Normal Correlation," *Biometrika* 31, no. 1/2 (1939): 9–12; W. A. Morgan, "A Test for the Significance of the Difference between the Two Variances in a Sample from a Normal Bivariate Population," *Biometrika* 31, no. 1/2 (1939): 13–19.

Table 4: The plan route and the Kelly route with learning on new business, side by side, means over 200 seeds; the Kelly route is the main construction. The prices are, for the Kelly route, its mean deviation from its own undeviated price on renewals, the share of its renewal quotes within one grid step of the band's upper end, and its mean deviation on new business, each deviation a logarithm pooled over every quote of the measured years and seeds and stated in percent as its exponential less one (§5.2); and, for each route, the premium written over the true expected loss of that business. The business written gives, for each route, the true expected loss of the business it writes over its own estimate. Table G.2 gives the deviations as the mean over seeds of each seed's mean. The rows on the path, the course of growth through the measured years, give the years in which the Kelly route trails by a resolved difference, and the Kelly route's deepest difference from the plan route, each read from the difference between the routes at the close of each measured year, paired by seed; the seeds in which it trails at the close; and each route's volatility, the standard deviation over the measured years of the yearly increment of growth, and its dispersion, the standard deviation across seeds of growth at the close, each with the t of the difference between the routes (for the dispersion, the Pitman–Morgan t of two paired standard deviations). Source: authors' calculations.

	The plan route	The Kelly route
The prices		
Renewals, percent from its undeviated price	—	+4.3
Renewal quotes within one grid step of the band's upper end, percent	—	14
New business, percent from its undeviated price	—	+1.2
Premium written over the true expected loss, new business	1.46	1.40
Premium written over the true expected loss, renewals	1.52	1.50
The business written		
True expected loss over the insurer's estimate, new business	0.989	0.989
True expected loss over the insurer's estimate, renewals	0.997	1.005
The path		
Measured years in which the Kelly route trails, by a resolved difference	—	1–28
Deepest difference from the plan route, percent of that year's surplus, in year 17	—	−10.4
Seeds in which the Kelly route trails at the close, of 200	—	173
Within-path volatility of growth, logarithm	0.0214	0.0195
Its difference, the Kelly route less the plan route: t	—	−7.45
Dispersion of growth across seeds at the close, logarithm	0.0926	0.1178
The ratio of the two: Pitman–Morgan t	—	+4.21

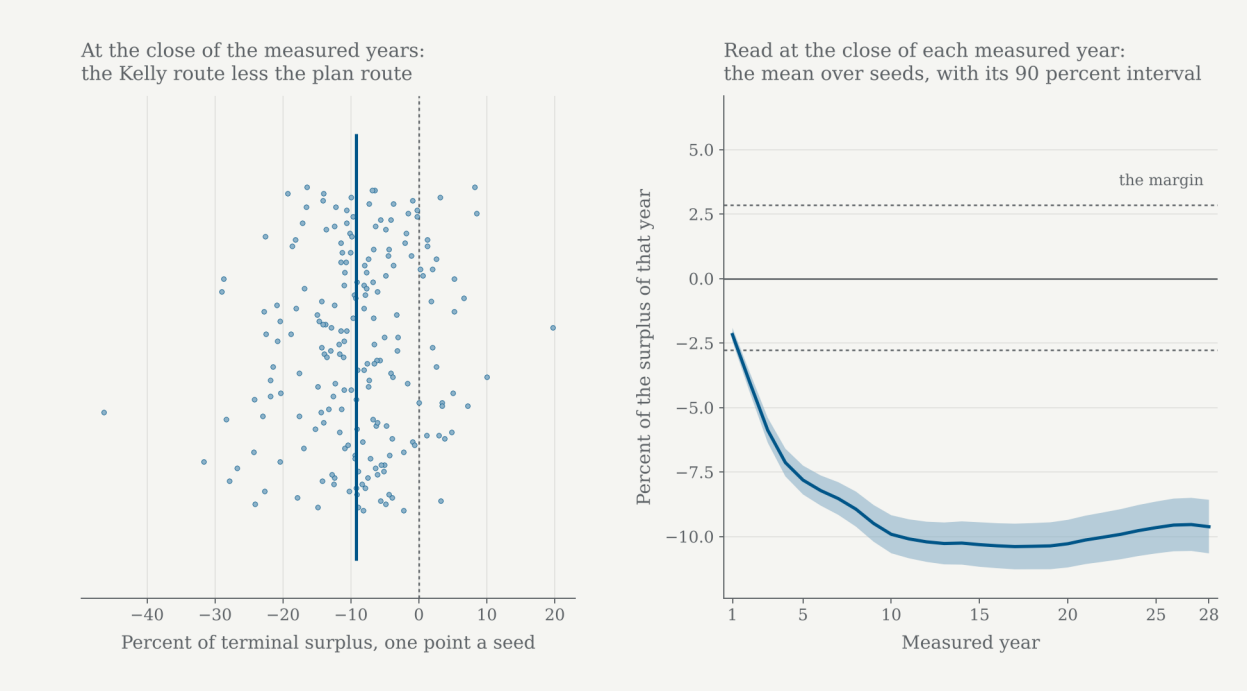


Figure 3: With learning on new business the Kelly route trails the plan route from the first measured year. Both panels give the Kelly route less the plan route. At left, the fourth seat's growth at the close of the measured years paired by seed, in percent of terminal surplus (horizontal axis), one jittered point a seed; the solid bar is the mean and the short-dashed line zero. At right, the same difference read at the close of each measured year (horizontal axis), in percent of the surplus of that year, the mean over seeds as the line and its 90 percent interval as the shaded area; the dashed lines are the equivalence margin either side of zero. The Kelly route is the main construction. Source: authors' calculations.

6 Other Settings and Other Markets

This section reports other settings of the main market and four further markets: the plan route with learning on new business in each, and the two routes where both were run. The main market is one setting of several dimensions an insurer does not choose. The plan route with learning on new business was run at other values of four of them, one at a time: the threshold at which the Bureau insurer refiles (0.05 and 0.15, against 0.10 in the main market); the noise of the underwriters' read of an account (0.35 and 1.05, against 0.7); the filed band, at 40 percent either way, against 25; and the buyer's responsiveness to price, at a scale of its response of 0.025 and 0.075, against 0.05 (a smaller scale is a more responsive buyer). It was also run in four further markets. In the first the insurers are smaller: each endowment is 0.467 of the main market's, at which an insurer's opening book stands at the capital ratio at which broker flight begins. The second has one tenth of the accounts, each endowment set by the main market's rule (§A.5) and so 0.100 of the main market's. In the third, the GLM insurer's seat is taken by a second insurer that learns and prices as the plan route on the learned estimate does. The fourth has a severity trend of 6.8 percent a year that no insurer projects, the yearly growth of the severity of other liability occurrence claims from 2015

through 2024,[11] applied to the mean claim limited per occurrence. The two routes, with learning on new business and on the learned estimate, were run at the wider band, at both values of the buyer's responsiveness, and in the four further markets. At the wider band the schedule modification and the schedule credit of every insurer move within 40 percent, and the Kelly route's hold curve is set at that band: its grid of deviations runs over the band (Appendix E), and the prior on the center of its hold curve has the band's spread (§D.2.2).

Table 5: The plan route with learning on new business at the close of the measured years, in the main market, at other values of four of its dimensions, and in four further markets: less the plan route on the learned estimate, and less the benchmark, in the same market. Each is a difference in terminal surplus in percent, paired by seed over 200 seeds, with its t statistic; a dash marks a setting at which the second comparison was not run. Source: authors' calculations.

	Less the plan route on the learned estimate		Less the benchmark	
The market	Percent	t	Percent	t
The main market	$+8.8$	$+12.48$	$+62.4$	$+38.92$
Refiling threshold 0.05 (lower)	$+8.6$	$+12.61$	$+61.0$	$+38.47$
Refiling threshold 0.15 (higher)	$+8.0$	$+9.75$	$+59.8$	$+36.73$
Read noise 0.35 (lower)	$+2.2$	$+2.52$	$+66.9$	$+38.05$
Read noise 1.05 (higher)	$+10.1$	$+14.65$	$+61.7$	$+40.19$
Filed band of 40 percent (wider)	$+3.0$	$+3.89$	—	—
Buyer's scale 0.025 (more responsive)	$+6.9$	$+8.34$	—	—
Buyer's scale 0.075 (less responsive)	$+7.9$	$+11.77$	—	—
Smaller insurers	$+14.0$	$+15.77$	$+78.9$	$+32.70$
One tenth of the accounts	-2.2	-1.29	$+24.5$	$+10.21$
A second insurer that learns	-8.4	-8.82	$+28.6$	$+21.74$
Severity trend of 6.8 percent a year	-7.2	-3.04	$+9.3$	$+3.24$

The plan route with learning on new business grows more than the plan route on the learned estimate at every other value run of the main market's dimensions and with smaller insurers (Table 5). The difference runs from 2.2 percent of terminal surplus, where the underwriters read an account with the least noise (not at the level that allows for the number of tests, §7.2), to 14.0 percent, where the insurers are smaller. In two of the further markets the order reverses. Beside a second insurer that learns, the plan route with learning on new business ends 8.4 percent below the plan route on the learned estimate, and under the severity trend 7.2 percent below (not at that level); the paper does not measure why. In the market of one tenth of the accounts the difference is not resolved (-2.2 percent, $t = -1.29$). Against the benchmark, the plan route with learning on new business grows more wherever the two were run, by 9.3 percent under the severity trend (not at that level) to 78.9 percent with smaller insurers.

[11] Jim Lynch and William Nibbelin, *Increasing Inflation on Liability Insurance – Impact as of Year-End 2024* (Insurance Information Institute and Casualty Actuarial Society, October 2025), 18; for other liability occurrence, net ultimate loss and defense and cost containment expense per claim, accident years 2015 through 2024.

Table 6: The Kelly route less the plan route at the close of the measured years, with learning on new business and on the learned estimate, in the main market and at each setting and in each further market where both routes were run with learning on new business; and the Kelly route with learning on new business less the Kelly route on the learned estimate. Each is a difference in terminal surplus in percent, paired by seed over 200 seeds, with its t statistic. The Kelly route is the main construction; at the wider band its hold curve is set at that band (Section 6). Source: authors' calculations.

	The Kelly route less the plan route				The Kelly route,	
	Learning on new business		The learned estimate		the first less the second	
The market	Percent	t	Percent	t	Percent	t
The main market	−9.6	−14.49	+1.1	+1.97	−2.7	−3.52
Filed band of 40 percent (wider)	−22.2	−40.22	−19.7	−42.06	−0.2	−0.35
Buyer's scale 0.025 (more responsive)	−11.1	−17.76	−3.4	−5.56	−1.6	−2.03
Buyer's scale 0.075 (less responsive)	−7.6	−13.74	+3.1	+5.51	−3.3	−4.47
Smaller insurers	−12.0	−17.60	+2.3	+3.19	−1.9	−1.93
One tenth of the accounts	−1.2	−0.73	+12.3	+8.50	−13.9	−7.85
A second insurer that learns	−3.5	−4.33	+0.3	+0.44	−11.9	−13.27
Severity trend of 6.8 percent a year	−4.6	−2.41	+7.2	+4.89	−17.4	−10.85

With learning on new business the Kelly route trails the plan route at every setting and in every market where both were run but the market of one tenth of the accounts (Table 6); under the severity trend, not at that level. The shortfall runs from 3.5 percent of terminal surplus beside a second insurer that learns to 22.2 percent at the wider band; in the main market it is 9.6 percent. In the market of one tenth of the accounts the difference is not resolved (−1.2 percent, $t = -0.73$). On the learned estimate the sign of the difference changes from setting to setting: the Kelly route trails at the wider band and with the more responsive buyer and leads with the less responsive buyer. In the further markets it leads with smaller insurers (not at that level), with one tenth of the accounts, and under the severity trend, and ends within the equivalence margin beside a second insurer that learns. There the 90 percent interval of the difference runs from −0.0092 to +0.0159 in the logarithm, inside the margin set for the main market, a margin set after this difference had also been measured. Learning on new business takes growth from the Kelly route at every setting but the wider band and in every market but that of smaller insurers, by 1.6 percent with the more responsive buyer to 17.4 percent under the severity trend; in the main market it takes 2.7 percent. The readings in the main market and with the more responsive buyer are not at that level. At the wider band and with smaller insurers, the difference is not resolved (−0.2 percent, $t = -0.35$, and −1.9 percent, $t = -1.93$). Appendix F states the predictions written for these runs before they ran and how each fared; the Kelly route on the learned estimate in the four further markets was run without a prediction. Appendix G gives the Kelly route's quotes with learning on new business by setting and market.

7 What Follows for an Insurer, and Conclusion

This section states what the results mean for an insurer's pricing, the limits of the measurement, and the conclusion.

7.1 Where Learning and the Kelly Decision Add Growth

In the main market, on the plan route, the learned estimate adds growth where the insurer holds no file of its own and takes growth away where it holds one. On new business the insurer quotes an account on the broker's loss run, and there the learned estimate adds 62.4 percent of terminal surplus over the rating plan's estimate. On a renewal the insurer holds the account's claim file, and there the rating plan's estimate, computed in part from that file through the experience modification, grows surplus more than the learned estimate: learning on renewals alone takes away 15.1 percent. The plan route with learning on new business reaches the highest growth of any insurer the paper reports in the fourth seat in the main market, 3.71 percent a year. It also grows more than the plan route on the learned estimate at every other value run of the main market's dimensions (at the lower read noise, not at the level that allows for the number of tests, §7.2) and with smaller insurers. In a market with a second insurer that learns, and in one with a severity trend that no insurer projects, the plan route on the learned estimate grows more (under the trend, not at that level); in the market of one tenth of the accounts the difference is not resolved (Section 6). In the plan route with learning on new business the rest of the rating plan is unchanged: the rate level, the schedule credit on new business, the provision for expenses and profit, and, on renewals, the experience modification. On new business the learned estimate takes the place of the class loss cost and the experience modification, with the schedule modification multiplying it as before.

The Kelly decision is worth most to an insurer that prices on its rating plan's estimate, an estimate on which the paper runs the Kelly route in the main market alone. On that estimate the Kelly route ends with 21.8 percent more terminal surplus than the plan route. On the learned estimate, in the main market, the Kelly route reaches the plan route's growth by a different book and a different path, within the equivalence margin. That margin was set after two differences between the routes had been measured: the one in the main market and the one beside a second insurer that learns. Across the other settings at which the Kelly route was run, the sign of its difference from the plan route changes. In the further markets the Kelly route leads (with smaller insurers, not at that level), except beside a second insurer that learns, where it ends within the margin (Section 6). With learning on new business the Kelly route falls short of the plan route, in the main market, at the other settings at which it was run, and in the further markets wherever the difference is resolved (under the severity trend, not at that level). In the main market, an insurer that prices on a learned estimate gets the gain from learning through its rating plan, without a decision computed at quote time; on the learned estimate, at one other setting and in three further markets (with smaller insurers, not at that level), the Kelly route adds to it (Table 6).

7.2 Limits

The main market is one line of business with independent traits, no loss trend (Section 6 reports a market with one), no shock common to all accounts, no reinsurance, no underwriting cycle.[12]

[12] On the explanations of underwriting cycles: Mary A. Weiss, "Underwriting Cycles: A Synthesis and Further Directions," *Journal of Insurance Issues* 30, no. 1 (2007): 31–45.

The three conventional insurers are built to price as the sources describe a rating plan, a loss-ratio indication, and a generalized linear model. Beyond the schedule credit and, for the Loss-ratio insurer, the complement of credibility set to the level of competitors' winning prices, they have no response to a competitor that learns. The market with a second insurer that learns gives the fourth seat a competitor that learns from the start, not one that responds to it. The buyer responds to price alone, through a rule of shopping and switching calibrated to one target, from a survey of 2007 of middle-market companies changing a property–casualty or employee-benefits insurer. The broker tells every insurer that lost at market its figure, rounded and without error; a figure that is partial, absent, or shaded is not run. Expenses are 0.30 of premium on new business and renewals alike, so a higher cost of acquiring an account is not measured. Surplus bears no income tax and pays no dividend, and terminal surplus carries each insurer's reserves at its own estimates, unrestated (§A.5). The schedule credits on new business and the Kelly route's deviations respond to competitors' prices. The Florida rule on schedule rating from which the main market's band is taken (Appendix A) excludes competitors' prices from what a schedule may rate. For personal lines, a task force of the National Association of Insurance Commissioners (NAIC) considers rate adjustments by a buyer's price elasticity of demand or propensity to shop inconsistent with the standard that rates not be unfairly discriminatory.[13] The model follows the practice the sources report and does not assess compliance.

The grid is run in the main market. The plan route, with learning on new business and on the learned estimate, is run at the settings and in the markets of Section 6, one dimension at a time, and the Kelly route, with learning on new business and on the learned estimate, at three of the settings and in the four further markets. The equivalence of the two routes' growth on the learned estimate in the main market, and beside a second insurer that learns, is a statement within a margin set with both differences already measured, at the precision of 200 seeds, and at the close of the measured years. What is learned and where it is used are each measured with the other part held at its class value or at the rating plan's estimate. Learning on new business divides learning by the kind of business and in no other way; a division by the length of an account's file or by its size is not run, and the paper does not measure why, in the main market, learning on renewals falls short of the rating plan's estimate beyond the selection of the renewals kept. At the main market's surplus the decision's variance term, the criterion's weight on risk, is equivalent to a charge of 0.0641 percent of the expected loss of an account at the bureau's scale of claim sizes. Of the main market and the four further markets, the variance term is largest in the market of one tenth of the accounts, a charge of 0.6421 percent of the expected loss of an account at the bureau's scale of claim sizes. For the Kelly route on the learned estimate, the variance term's removal is run in two markets (§D.4). In the market of one tenth of the accounts, removing the term lowers the route's growth by 2.3 percent of terminal surplus, a difference of -0.023 in the logarithm (not at the level that allows for the number of tests, below). In the main market, removing it is shown not to move the route's growth by more than 0.010 in the logarithm. There the paper therefore tests the criterion's rule for pricing under uncertainty about the odds: a margin weighed against the chance of winning the account and the discounted value of keeping it. The property that separates the criterion from the maximization of expected profit, the weight the logarithm puts on the variance of an outcome against surplus, shows when the stake is large against the surplus, and the paper does not run a market whose accounts are large against an insurer's surplus. The paper runs the full criterion and not a fractional criterion, which scales the criterion's stake by a constant below one, the form the literature uses to give up growth for a lower chance of a large loss of surplus. An insurer that

[13] National Association of Insurance Commissioners, Casualty Actuarial and Statistical (C) Task Force, *Price Optimization White Paper* (adopted by the task force, November 19, 2015), 15–16.

weighs a loss of surplus differently from the logarithm need not prefer the criterion.[14] The anchored lifetime value is set at the undeviated price, which is below the Kelly route's own renewal prices on average, so the anchored lifetime value understates the lifetime value. At any quote, other inputs held, the decision's deviation is therefore higher than it would be with a lifetime value set at the Kelly route's own renewal prices; the paper does not measure by how much (§D.1.5). The horizon is known to the insurer, the value of keeping an account falls toward zero over the final measured years, because its sum stops at the horizon, and growth is read at that horizon. In the main market the difference between the two routes on the learned estimate narrows after its deepest, in year 10, and its 90 percent interval lies inside the margin from year 22 through the close; the paper does not divide that narrowing between the decision and the approach of the horizon, or trace it to any feature of the two constructions. It does not establish the cause of the Kelly route's shortfall with learning on new business. And it does not run a market whose horizon comes earlier, or one under a going-concern horizon. The paper does not separate an insurer's growth into what the size of its book contributes and what its margin contributes. It does not measure why the order of the two estimates on the plan route reverses beside a second insurer that learns and under the severity trend (Section 6), or why the new business the plan route on the learned estimate wins carries a true expected loss of 0.927 times its estimate (§4.1). That a markup today does not change next year's margin is argued from the market's rules and not measured (§2.1, §D.1.5). No insurer fails in any seed, so the decision's exposure to ruin is not tested (§A.5). And every difference is a total effect within a market that responds through its competitors' capital and books, measured over 200 seeds of one model, so that its t statistic is a statement about the simulation's precision at that count and not about an observed market. Each prediction is read by the rule registered for it, without adjustment for the number of predictions tested. At the level that allows for the number of tests, $|t| > 3.68$ (Appendix F), three readings on which Section 6 and this section draw do not hold: learning on new business against the plan route on the learned estimate at the read noise of 0.35, and, under the severity trend, learning on new business against the benchmark and, with learning on new business, the Kelly route against the plan route. Nor, at that level, do five further readings these sections draw on: the reversal of the two estimates on the plan route under the severity trend (Table 5); in Table 6, the Kelly route's lead on the learned estimate with smaller insurers and the growth that learning on new business takes from the Kelly route in the main market and with the more responsive buyer; and the growth the Kelly route on the learned estimate loses without the variance term in the market of one tenth of the accounts. Appendix F names one further prediction that held under the registered rule and not at that level, the seventh registration's sixth, on which these sections do not draw.

7.3 Conclusion

A rating plan gives an insurer a price. Under the Kelly criterion the price is for the growth of the insurer's surplus, which, in repeated independent bets on known odds, the criterion maximizes in the long run. Beyond the expected loss a plan also estimates, the criterion requires three things a plan does not hold: the variability of the loss against surplus, the value of keeping the account, and the chance that the quote wins. An insurer can learn all four from the records of its own operations, with what the market makes public. This paper has measured, in the main market and in part at

[14] Paul A. Samuelson, "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long," *Journal of Banking and Finance* 3, no. 4 (1979): 305–307; L. C. MacLean, W. T. Ziemba, and G. Blazenko, "Growth versus Security in Dynamic Investment Analysis," *Management Science* 38, no. 11 (1992): 1562–1585; Leonard C. MacLean, Edward O. Thorp, and William T. Ziemba, eds., *The Kelly Capital Growth Investment Criterion: Theory and Practice* (Singapore: World Scientific, 2011).

other settings and in four further markets, over 200 seeds, what learning the loss is worth, where it is worth it, and what the Kelly decision, with the hold curve it learns and the lifetime value it reads from its own book, adds on each estimate.

In the main market learning the loss is the largest single effect, and learning an account's claim frequency is worth more than learning the size of its claims. On the plan route the worth of learning the loss lies in new business, the accounts an insurer does not yet hold; on the accounts it holds, the rating plan's estimate, computed in part from the insurer's own claim file through the experience modification, grows surplus more. The plan route with learning on new business grows its surplus at 3.71 percent a year, against 1.93 percent for the benchmark in its place, which does not learn the loss. In the main market it ends 8.8 percent of terminal surplus above the plan route on the learned estimate. The plan route with learning on new business keeps that lead at every other value run of the main market's dimensions (at the lower read noise, not at the level that allows for the number of tests, §7.2) and with smaller insurers, and loses it beside a second insurer that learns and under a severity trend that no insurer projects (under the trend, not at that level). In the market of one tenth of the accounts the difference is not resolved. In the main market the Kelly route grows more than the plan route on the rating plan's estimate, ends within the equivalence margin of the plan route on the learned estimate, a margin set after the difference was measured, and falls short of the plan route with learning on new business: in the main market, the better the estimate, measured by the growth it gives at the plan's pricing, the less the Kelly decision adds. In the two markets where the plan route with learning on new business loses its lead (under the trend, not at that level), the Kelly route on the learned estimate ends above the plan route on the learned estimate under the severity trend, and within the equivalence margin of that route beside a second insurer that learns, a margin set after that difference too was measured. Where both routes learn on new business, the Kelly route trails the plan route in both markets (under the severity trend, not at that level; Table 6).

Any estimator of each account's loss distribution and of the chance of winning it fits the paper's grid: it can be used on new business or on renewals and priced by the plan route or by the Kelly route, and the paper's measure of what learning is worth is the measure such an estimator would be held to. The learning of this paper uses standard statistical methods: credibility weighting at moments estimated from the data an insurer holds, and a hold curve fitted to the quotes it won and lost. What the criterion's weight on risk would add in a market whose accounts are large against an insurer's surplus, how the routes fare where all three conventional insurers learn, and what a neural network fitted to the same data would add, free of the stated form these estimators take for the relations among an account's characteristics, its claims, and the market's responses, are questions this paper leaves open.[15]

[15]A network with one hidden layer and enough units can approximate any continuous relation on a bounded range of inputs to any stated accuracy: Kurt Hornik, Maxwell Stinchcombe, and Halbert White, "Multilayer Feedforward Networks Are Universal Approximators," *Neural Networks* 2, no. 5 (1989): 359–366.

A The Market: Its Assumptions, Its Insurers, and Its Loss Process

This appendix states the market of Section 3 in full: its assumptions (Table A.1), the accounts and their losses, what an insurer knows, the three conventional insurers and how each prices, the buyer and the broker, capital and the measure of growth, and the loss process.

A.1 Accounts and Losses

The market holds 4,000 accounts of one line, written by four insurers without reinsurance. An account has a size, lognormal with median 1 and a log standard deviation of 1, and two traits, one of claim frequency and one of claim size, independent of each other. Each trait, in logarithms, is the sum of a permanent part and a part that drifts as a first-order autoregression with persistence $\rho = 0.85$: the frequency trait's permanent part has a standard deviation of 0.40 and its yearly shock one of 0.30, and the severity trait's permanent part one of 0.632 (a variance of 0.40) and the same shock. In a policy year the account has a Poisson number of claims with mean λ, the product of its size, its frequency factor (§A.6), and one expected claim a year at the median size. A claim's size is mixed exponential, the form used in increased-limits ratemaking, with means of \$9,277, \$92,766, \$927,662, and \$9,276,615 at weights of 0.700, 0.250, 0.045, and 0.005, scaled by the account's severity factor (§A.6); the means are set, through one multiplier of 1.85532 on a stated base curve, so that the book's mean claim limited at the occurrence limit is Verisk's general liability indemnity severity of \$56,213 a claim.[16] Indemnity is limited to \$1,000,000 an occurrence and \$2,000,000 in the aggregate of a policy year, and loss adjustment expense is 0.39156 of the indemnity paid, in addition to the limit, the ratio of Verisk's loss adjustment expense severity of \$22,011 to its indemnity severity. A claim is reported with a lag, 0.55 of a policy year's claims by the end of the year and all of them within four years, and is paid once, at settlement, up to 12 years after it is reported. Until it settles it stands on an insurer's books at a case reserve, its ultimate value times a lognormal error. §A.6 states the process in full.

A.2 What an Insurer Knows

Three sources of information exist, and they are not equally shared. The first is public. A classification score, the sum of an account's two permanent traits read with noise, places each account in one of five classes, and a bureau publishes for each class the market's pooled experience of five policy years ending two years before the rating year: a loss cost, a loss cost with each occurrence and its expense capped at \$100,000, and a claim frequency, each per unit of size. The market also treats as public the bureau's curve of claim sizes and its moments at scale one, the reporting and settlement patterns, the spread of the case-reserve error on a loss run and on a holder's file, the share of loss adjustment expense, and the distribution of the traits within a class (§D.2). The second is the submission. An insurer that is asked to quote receives the broker's loss run, the account's claims of the past five policy years whoever wrote them, valued one quarter before the policy's inception,[17] with an open claim at a case reserve whose error has a log standard deviation of 0.30. Each insurer's

[16]The form: Clive L. Keatinge, "Modeling Losses with the Mixed Exponential Distribution," *Proceedings of the Casualty Actuarial Society* 86 (1999): 654–698; Joseph M. Palmer, *Increased Limits Ratemaking for Liability Insurance* (Casualty Actuarial Society study note, 2006). The severity: Verisk, *2024 Executive Insights: General Liability*, 1st ed., Verisk's Core Lines Services Report, Second-Quarter 2019 to First-Quarter 2024 (Verisk, 2024), 4.

[17]Association for Cooperative Operations Research and Development, *Commercial Insurance Application*, ACORD 125 (2007/10), 3, which asks for the claims or losses of the prior five years.

underwriter also forms a read of the account of its own, the permanent traits read with noise of the size of the score's, a standard deviation 0.7 times that of the traits. The third is private. The insurer that holds an account reads its own claim file for every year it has held it, current to the quarter and with half the case-reserve error, and keeps the loss run it received when it acquired the account. The holder therefore knows more about its accounts than a challenger does, and the longer it has held an account the more: this is the asymmetry that the theory of insurance pricing under asymmetric learning studies (Section 1), and nothing else in the market favors a holder except the buyer's costs of shopping and switching.

A.3 Three Conventional Insurers

Each conventional insurer prices by the bureau's plan and differs from the others in how it sets its level. For an account of class c and size z the plan's expected loss is

$$\bar{x} = L_c\, z \times M_{\text{exp}} \times M_{\text{sch}} \times R, \tag{A.1}$$

where L_c is the bureau's loss cost, M_{exp} the experience modification, M_{sch} the schedule modification, and R the insurer's rate level, and the premium is $\bar{x}/(1-\pi)$ with a provision π = 0.34 for expenses of ε = 0.30 of premium and a profit of 0.04.[18] The experience modification weights the account's capped losses against the class's by credibility: $M_{\text{exp}} = 1 + Z_{\text{exp}}(A - E_{\text{rep}})/E$, with A the account's losses on the file the insurer reads, each occurrence with its expense capped, E the class's expected capped losses for the period, E_{rep} their expected reported part, and $Z_{\text{exp}} = E/(E+K)$. The bureau's plan as filed uses three policy years, the latest year before the rating year left out, and K of \$66,000: the plan states K as \$100,000 of premium, and it enters on the expected-loss basis, times $1-\pi$, so that a change of the provision moves K with it. An insurer applies the plan's formula over every year its file holds, the latest again left out, a year of age j weighted by d^{j}, and estimates d and K from its own book by least squares (Appendix B). The schedule modification is what the underwriter's own reading adds to the class, within the filed band, 25 percent either way in the main market and 40 percent at the wider band of Section 6.[19]

The *Bureau insurer* files a rate level, a multiplier on the bureau's loss costs, and refiles it when the indication on its own experience reaches the threshold below. At each year's close it computes the loss-ratio method's indication on its own experience of the past five policy years: with a its reported losses, e the expected reported losses it priced those years at, brought to its present level, and n its reported claims, the indicated change is

$$Z f + (1 - Z)\,\psi, \qquad f = a/e, \qquad Z = \min\!\big(1, \sqrt{n/1{,}082}\big), \tag{A.2}$$

with the complement of credibility ψ equal to one.[20] The standard of 1,082 claims is the full-credibility standard of classical (limited-fluctuation) credibility for claim frequency, applied here

[18] National Association of Insurance Commissioners, *Report on Profitability by Line by State in 2024* (Kansas City, MO: NAIC, 2026), 9: for other liability countrywide, a line wider than general liability, general expense, selling expense, and taxes, licenses, and fees of 8.2, 18.8, and 1.9 percent of net premiums earned, 28.9 percent in all.

[19] Fisher et al., *Individual Risk Rating*; a limit of 25 percent either way is that of Florida Administrative Code, Rule 69O-170.004. The same rule, at (2)(b), excludes from the characteristics a schedule may rate "the degree of competition for the risk or the rates which may be offered by other insurers"; the schedule credit on new business described below and the Kelly route's deviation within the band (§D.1.4) answer competitors' prices, as underwriters use schedule rating to meet a competitor's quote in practice, and the model follows the practice: Lisa A. Hays, "Managing Commercial Lines Pricing Levels in a Loss Cost Environment," *Casualty Actuarial Society Forum* (Winter 2002): 11.

[20] Laurence H. Longley-Cook, "An Introduction to Credibility Theory," *Proceedings of the Casualty Actuarial Society* 49 (1962): 194–221; Werner and Modlin, *Basic Ratemaking*, chaps. 8 and 12, and apps. A–D, the worked indications.

to losses, as practice often applies it. Where the indicated change differs from one by 0.10 or more it refiles, and its level is multiplied by the change; otherwise its level stands.

The *Loss-ratio insurer* moves its level every year by Equation A.2 on a window of five policy years, the span of Werner and Modlin's worked indications, with the complement ψ set to competitors' rates: the level of the winning prices relative to its own plan prices, fitted as a censored normal to the broker's figures on the quotes it lost at market (§A.4) and to the quotes it bound (Appendix B). The complement takes that level once the window holds 100 lost quotes, and one before.

The *GLM insurer* adds relativities of its own to the plan. Each year it fits a multiplicative model with a logarithmic link to the capped reported losses of its own book over five policy years, with a factor for each class and for each of three size bands, thirds of the year's accounts by manual premium, the premium at the bureau's loss cost for the account's class and size before the experience and schedule modifications, and takes each fitted factor at its classical credibility; its level moves by the ratio of its book's reported losses to their expected amount at its present relativities, taken at the classical credibility of the book's claims (Appendix B).[21]

On new business, a quote on an account it does not hold, a conventional insurer files no load. In every one of Wu and Lin's books new business runs at a higher loss ratio than renewals. In their two books of general liability, new policies receive schedule debits about half as often as renewals and schedule credits as often: there, new business is priced in practice no higher than the plan states, and the loss-ratio difference is a cost the price does not carry. Wu and Lin also judge a surcharge on new business justifiable on their data, so a filed load of zero follows the practice their general liability books show and not their reading of what the data justify.[22] What an underwriter may move without a filing is the schedule modification, and insurers do not always implement the rate their experience indicates.[23] In the middle and the largest band the insurer therefore applies a schedule credit or debit to its new-business quotes, equal to the level of the winning prices it has seen in the band relative to their level in the smallest band, from the broker's figures of the past three years, with the schedule credit and the schedule modification together kept inside the filed band. The smallest band is the reference and carries no schedule credit. The rule refers to account size only through the choice of the smallest band as the reference.

A.4 The Buyer and the Broker

An account renews once a year, in the quarter of its inception. Its holder quotes. The broker invites one challenger, drawn at random among the other insurers that are open to new business, and each further one with a probability equal to the account's rank by manual premium among the year's accounts, to the power 1 (Table A.1), from zero for the smallest through one for the largest, so that the smallest account draws one competing quote and the largest draws all three. Let Δ be the logarithm of the holder's quote over the lowest invited quote. With probability 0.10 the account is remarketed for reasons unrelated to price and chooses among all the quotes by a logit on log price with the scale s of the response below. Otherwise it goes to market with probability $1/(1+\exp(-(\Delta-k)/s))$ and, once there, stays with its holder with probability $1/(1+\exp((\Delta-k')/s))$, taking the lowest invited quote if it does not; k is the cost of shopping, k' the cost of switching, and s = 0.05 the scale of the response, each in units of log price. The two costs are equal and calibrated

[21] Robert A. Bailey, "Insurance Rates with Minimum Bias," *Proceedings of the Casualty Actuarial Society* 50 (1963): 4–11; Werner and Modlin, *Basic Ratemaking*, chap. 10.

[22] Wu and Lin, "Large Scale Analysis," 398, 402, and 405.

[23] Werner and Modlin, *Basic Ratemaking*, chap. 13; Hays, "Commercial Lines Pricing Levels," 11.

rather than stated. On other seeds, $k = k'$ was moved until 0.14 of the accounts of the market with the benchmark in the fourth seat changed insurer a year, half the 28 percent of middle-market companies that a 2007 survey by Greenwich Associates found changing a property–casualty or employee-benefits insurer over two years; the halving counts no company changing twice, and independent yearly switching would give 0.151. The calibration gives $k = k'$ = 0.061184. At that cost, in the calibration's own run, 0.1404 of the accounts change insurer a year (standard error over 20 seeds 0.0033), 0.2654 go to market, and 0.8313 of premium is retained, between two retentions of commercial business published in 2026.[24]

After placement, each insurer that quoted on an account that went to market and lost is told by the broker the difference between its quote and the bound premium, as a fraction of the bound premium rounded to the nearest 0.05. The insurer that bound is told nothing, and nothing is said of an account that did not go to market. This is the only information about competitors' prices that any insurer receives.

A.5 Capital, Solvency, and the Measure

Within a seed every insurer starts the first measured year with the same endowment, set so that a seat's expected written premium of the opening book and the steady-state reserves of a book of that premium stand at the industry's aggregate risk-based-capital (RBC) ratio of 6.42458 (Appendix C); over the 200 seeds it averages $556.7 million, from $519.3 through $605.8 million.[25] A policy's premium is booked with expenses of 0.30 and a reserve at the writer's own expected loss; the reserve runs off as claims are reported and paid; and surplus and reserves alike earn the *return on funds*, r = 0.0356 a year, the industry's published investment yield.[26] At each year's close an RBC ratio is computed for each insurer, its surplus over a requirement of the NAIC's form in its written premium and its reserves (Appendix C). Capital affects behavior at three thresholds. The broker reads the ratio: below 3 it withholds a new-business invitation with a probability that rises to one at 2 (broker flight), and a renewal is never affected. The insurer reads its own writings: in a quarter that opens with its written premium of the year above three times its surplus, the level at which the NAIC's Insurance Regulatory Information System (IRIS) ratio of premium to surplus is flagged, it quotes no new business, and it keeps quoting renewals. And below a ratio of 0.70, or with its surplus exhausted, the insurer fails and is replaced by an entrant, a Bureau insurer with the endowment, at most two times in a seed.[27]

The books are built over 19 years in which accounts change holder at random at a rate of 0.15 a year, so that every insurer opens with a seasoned book and a file on each account it holds. The

[24] Greenwich Associates, *Insurance Market Pulse: Looking at the Small Business and Middle Market Segments* (2007), the share of middle-market companies that changed a property–casualty or employee-benefits insurer over the two years before the survey, as reported in Douglas McLeod, "Changing Insurers Yields Surprise," *Business Insurance*, June 17, 2007. The retentions: The Hartford Insurance Group, Inc., *Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2025* (filed with the U.S. Securities and Exchange Commission, February 20, 2026), a middle-market premium retention of 83 percent; The Travelers Companies, Inc., "Travelers Reports Excellent Second Quarter and Year-to-Date Results," press release, July 17, 2026, a retention of 86 percent in its Business Insurance segment, retention being the premium available for renewal that was retained, excluding rate and exposure changes, with the segment's National Accounts left out.

[25] National Association of Insurance Commissioners, "Aggregated P&C RBC Data: 2025 Data as of June 1, 2026," *Summary: Aggregate P/C RBC Results by Year* (NAIC, 2026).

[26] National Association of Insurance Commissioners, *U.S. Property & Casualty and Title Insurance Industries—2025 Full Year Results* (NAIC Financial Regulatory Services, 2026).

[27] The failure threshold is the level at which the NAIC's model act mandates control: National Association of Insurance Commissioners, *Risk-Based Capital (RBC) for Insurers Model Act*, Model 312 (2012).

market then runs 10 practice years, in which every surplus is reset to the endowment at the year's close while rate levels, estimators, and books carry forward, and 28 measured years; at the close of the final practice year, before the reset, each insurer's reserves on the business it wrote before the measured years are restated, the provision for unreported claims at its true expectation and each open claim's case reserve at what the claim will pay (Appendix C). An insurer's *growth* is

$$g = \ln \frac{S_T}{S_0} - T \ln(1 + r), \qquad T = 28, \quad r = 0.0356, \tag{A.3}$$

the terminal logarithm of its surplus over the endowment, net of the return on funds, so that g is zero for a surplus that earns that return and nothing else, and g/T, a logarithm, is a yearly rate of growth from writing business, the investment income on the reserves an insurer holds included; its exponential less one is the rate a year the tables state. Surplus bears no income tax and pays no dividend. The terminal surplus S_T is taken from the insurer's own books at the close of the final measured year, its reserves at its own estimates: no restatement like that of the final practice year is made there, and the paper does not measure the difference one would make. Every *arm*, one insurer in the fourth seat in the main market or in a variation of it, runs on seeds 0 through 199, and every run is a function of its arm and its seed alone. An insurer that fails leaves the market, and its growth is not defined; no insurer fails in any seed of any arm, so every mean and contrast is over all 200 seeds, and the paper does not test the decision's exposure to ruin. A contrast between two arms or two insurers is paired by seed, with its mean, its t statistic (the mean over its standard error, the sample standard deviation over seeds divided by the square root of the number of seeds), and the number of seeds in which it is above zero. A contrast between two arms is a total effect (Section 3), not the effect of the change with the competitors held as they were.

A.6 The Loss Process

An account i has a size z_i, lognormal with median 1 and log standard deviation 1. Its frequency trait in year t is $u_{it} = \alpha_i + d_{it}$, with α_i normal with mean zero and standard deviation 0.40 and $d_{it} = \rho\, d_{i,t-1} + \eta_{it}$, ρ = 0.85, η_{it} normal with standard deviation 0.30, started from its stationary distribution. Its severity trait is of the same form, with a permanent standard deviation of 0.632 (a variance of 0.40) and the same persistence and shock, independent of the frequency trait. The stationary variance of a trait's drifting part is $0.30^2/(1 - \rho^2)$, so the frequency trait has a variance of 0.484 and the severity trait one of 0.724. Each trait enters the loss through its *factor*, the exponential of the trait less half the trait's variance, which has a mean of one across accounts in every year.

In policy year t the account's claim count is Poisson with mean

$$\lambda_{it} = \lambda_0\, z_i\, \exp\!\big(u_{it} - \tfrac{1}{2}\mathrm{Var}(u)\big), \qquad \lambda_0 = 1. \tag{A.4}$$

A claim's size before any limit is exponential with a mean drawn from the four means of §A.1 at their weights and multiplied by the account's severity factor; the curve's mean before any limit is \$117,813, and limited to \$1,000,000 an occurrence it is \$61,964 at scale one, against the book's mean of \$56,213 across the accounts' severity factors (§A.1). Within a policy year the aggregate limit of \$2,000,000 erodes in the order in which the claims are paid. Loss adjustment expense is 0.39156 of the indemnity paid and is outside the limits. An account's true expected loss in a year, the quantity the tables call by that name, is λ_{it} times the expected claim limited per occurrence at the account's severity factor, with loss adjustment expense, times the share of the year's expected occurrence-limited losses that the aggregate limit leaves.

Table A.1: The market's assumptions, each with its value in the main market and its grounding. An assumption is grounded in a source, a calibration to a sourced target, a source's form, the objective, or the simulation's scale, or it is stated, where no source was found and the value is the model's own. The settings of Table F.1 and Section 6 vary the Bureau insurer's threshold and the read noise; those of Section 6 also vary the filed band and the buyer's scale, and its further markets vary the accounts, the endowment, the severity trend, and the insurer in the GLM insurer's seat; at the wider band the Kelly route's hold curve is set at that band. The insurers of Table 1 vary the lifetime value's discount. Amounts are per occurrence or per policy year and rates per year; the sources are in the text. Source: authors' construction.

Assumption	Value	Grounding
Account size, log standard deviation	1	stated
Claims a year at the median size	1	stated
Frequency trait: spread, persistence, shock	0.40, 0.85, 0.30	stated
Severity trait: variance, persistence, shock	0.40, 0.85, 0.30	stated
Claim size, at Verisk's severity	mixed exponential	Keatinge; Palmer; Verisk
Loss adjustment expense over indemnity	0.39156	Verisk
Limits: occurrence; aggregate ($)	1,000,000; 2,000,000	Palmer
Share of claims reported by year	0.55, 0.85, 0.95, 1.00	stated
Settled within one, two, five years of report; all within	0.30, 0.50, 0.82; 12	stated
Case-reserve error: loss run; holder's file	0.30; 0.15	stated
Classes; size bands	5; 3	scale
Read noise, over the traits' spread	0.7	stated
Bureau's statement: years; lag	5; 2	stated
Loss run: years; valuation lag, quarters	5; 1	ACORD; stated
Full credibility, claims	1,082	Longley-Cook, frequency
Filed experience plan: years; cap; K ($)	3; 100,000; 66,000	bureau plan's form
Filed band, either way	0.25	Florida rule
Provision π; expenses ε; profit	0.34; 0.30; 0.04	NAIC; stated
Filed load on new business	0.00	Wu and Lin, GL books
The load's rule	schedule credit by band	stated
Lost quotes to fit: complement; schedule	100; 30	stated
Windows: loss-ratio method; GLM fit, years	5; 5	Werner and Modlin; stated
Lifetime value's estimation window, years	3	stated
Schedule fit's window, years	3	stated
Bureau insurer: experience, years; threshold	5; 0.10	stated
Costs of shopping and switching, $k = k'$	0.061184	calibrated to Greenwich
Scale of the buyer's response s	0.05	stated
Remarketing unrelated to price	0.10	stated
Invitations: power of the account's rank	1	stated
Broker's figure rounded to	0.05	stated
Capital requirement: ρ; κ	1.72456; 0.0968742	NAIC 2025
Flight band, RBC ratio (NAIC trend test)	2 through 3	NAIC; slope stated
New-business stop, premium over surplus	3	NAIC IRIS; the stop stated
Failure, RBC ratio	0.70	NAIC Model Act
Return on funds, on surplus and reserves	0.0356	NAIC yield
Endowment, RBC ratio	6.42458	NAIC 2025
Entrants after failures, at most	2	stated
Hold curve: prior spreads; decay	0.25, 1.0; 0.98	stated
Lifetime value: discount; summed through	return on funds; horizon	the construction's choice, both discounts run (Table 1); the objective
Accounts; years of build-up, practice, measured	4,000; 19, 10, 28	scale
Build-up: holders changing a year	0.15	stated

A claim occurs in a quarter of the policy year drawn uniformly. It is reported within the policy year with probability 0.55, by the end of the second year with probability 0.85, of the third with 0.95, and of the fourth with certainty: in the policy year, in a quarter drawn uniformly from its occurrence quarter through the year's final quarter, and in a later year uniformly over that year's quarters. It settles between one quarter and 12 years after it is reported, 0.30 of claims within a year of report, 0.50 within two, 0.82 within five, and all within 12. An open claim is carried at a case reserve, its ultimate value times $\exp(\sigma_c \xi - \sigma_c^2/2)$ with ξ standard normal, one draw a claim, and σ_c of 0.30 on a broker's loss run and 0.15 on the holder's own file. A settled claim stands at what was paid.

B The Rating Plans

This appendix states the rating plans in their formulas: the classes and the bureau's statement, the experience and schedule modifications, the GLM insurer's fit, and the schedule credit on new business.

The classification score is α_i of the frequency trait plus the permanent part of the severity trait plus normal noise with a standard deviation 0.7 times that of the sum, and the classes are the quintiles of its distribution. For policies of year t the bureau's statement covers policy years $t-6$ through $t-2$, valued at the start of year t. For each class it gives the reported losses, the reported losses with each occurrence and its expense capped at \$100,000, and the reported claims, each divided by the public expected reported share of its year and summed over the market, per unit of size; a class's figure is weighted with the all-class figure at the classical credibility of the class's reported claims.

In the experience modification, for an account of class c and size z let $B_c z$ be the expected capped loss of a year. Over the years of the file the insurer reads, a year of age j (the latest year of the period at age zero, the period ending two years before the rating year) weighted by $w_j = d^j$,

$$M_{\text{exp}} = 1 + \frac{E}{E+K} \cdot \frac{A - E_{\text{rep}}}{E}, \quad A = \sum_j w_j A_j, \quad E = B_c z \sum_j w_j, \quad E_{\text{rep}} = B_c z \sum_j w_j F_j, \tag{B.1}$$

with A_j the year's capped reported losses on the file and F_j the public expected reported share of the year at the file's valuation. The holder applies it to its own claim file and a challenger to the broker's loss run. Each insurer estimates (d, K) once a year by least squares over a grid, d from 0.5 through 1.0 and K from \$25,000 through \$500,000: for the accounts it held and renewed in the latest expired policy year it compares the year's expected capped loss times the modification each pair would have given, from the file as it stood when the year was quoted, with the year's capped losses as now reported, and it adds the squared differences, each divided by the year's expected capped loss, to those of the years before. Until its first estimate it uses the plan as filed, d = 1 over three years and K of \$66,000.

In the schedule modification, with (ν_1, s_1^2) the mean and variance of the permanent traits' sum in the account's class and (ν_2, s_2^2) its mean and variance given the class and the underwriter's reading, each under a normal approximation, the modification is $\exp(\nu_2 - \nu_1 + (s_2^2 - s_1^2)/2)$, held within 1 ± 0.25 in the main market and 1 ± 0.40 at the wider band of Section 6.

On the policies the GLM insurer wrote in the five policy years ending at the year's close, it fits $A \approx E \times f_{\text{class}} \times f_{\text{band}} \times f_{\text{new}}$, with A a policy's capped reported losses, E their expected reported amount at its present relativities, and f_{new} a factor by band on new business, by iterative proportional fitting, the estimator of a Poisson likelihood with a logarithmic link. The band factors and

then the class factors are brought to a mean of one, weighted by E. The change of level is the ratio of the book's reported losses over those years, uncapped and with expense, to their expected reported amount at the present relativities. The level, each class relativity, and each band relativity are then multiplied by $Zf + 1 - Z$, with f the change of level or the fitted factor and Z the classical credibility of the claims behind it. The factors on new business do not enter its price.

For the schedule credit on new business, the insurer takes for each band, from its quotes as a challenger of the past three years on accounts that went to market and were placed, the logarithm of the lowest competing premium over its own plan price at its present level: observed as the bound premium where the broker gave a figure, and known only to be at or above its quote where its own quote bound. The bound premium is the lowest competing premium unless the holder kept the account at a price above another challenger's quote. A normal distribution is fitted by maximum likelihood to both kinds of observation, and the band's level of winning prices, the logarithm of the expected ratio of the lowest competing premium to the plan price under the fitted distribution, is the fitted mean plus half the fitted variance. The schedule credit in the middle and the largest band is the exponential of the band's level less the smallest band's. A band's schedule credit changes only where the band and the smallest band each hold 30 lost quotes in the window; otherwise it keeps its previous one. The same fit over all bands, on its own window of five years, gives the Loss-ratio insurer's complement of credibility ψ once it has 100 lost quotes in the window.

C The Capital Standard

This appendix states the capital standard in its formulas: the risk-based-capital ratio, broker flight, the stop on new business, and the endowment the standard sets.

At a year's close an insurer with written premium W in the year, reserves Q, and surplus S has the risk-based-capital ratio

$$\mathrm{RBC} = \frac{S}{\kappa\sqrt{W^2 + (\rho\,Q)^2}}, \qquad \rho = 1.72456, \quad \kappa = 0.0968742. \tag{C.1}$$

The square-root combination is the form of the NAIC's formula, in which the premium-risk and reserve-risk charges combine under a root; the model carries those two charges and absorbs the others in κ in the industry's proportion, the covariance term of the two underwriting charges alone standing at 0.470 of the industry's total risk-based capital after covariance, of which the Authorized Control Level is half. The two constants are calibrated on one year of the NAIC's aggregated data: ρ is the reserve charge per dollar of reserves over the premium charge per dollar of net written premium, and κ is set so that the industry's own point—net premium at 0.773 of surplus, reserves at 1.06 times premium (the calibration uses the unrounded ratio)—stands at the industry's aggregate ratio of 6.42458.[28] A failed insurer's seat (§A.5) passes, while entries remain, to an entrant at a level of one with the endowment and no book. The model insurer's reserves run above the industry's all-lines ratio to premium, as a long-tail book's do, so at an equal ratio of premium to surplus its RBC ratio is lower than the industry's, as the formula's reserve charge makes it.

Broker flight is the broker's response to the ratio. At each year's close the broker reads each insurer's RBC ratio and, in the following year, extends a new-business invitation to it with probability $\min(\max(\mathrm{RBC} - 2, 0), 1)$: one at a ratio of 3 or more, falling linearly to zero at 2, account by

[28] National Association of Insurance Commissioners, "Aggregated P&C RBC Data"; National Association of Insurance Commissioners, *2025 Full Year Results*.

account by its own random stream. The band is the NAIC trend-test band and stands in the model as a proxy for the financial-strength thresholds brokers act on; the steepness of the withdrawal has no source. Flight is suspended in the practice years, and the first measured year opens on the final practice year's closing statement with the surplus reset to the endowment.

The stop on new business (§A.5) is set on IRIS ratio 2, net premiums written over policyholders' surplus, at the level the NAIC flags, on the year's written premium to date; the test is made again each quarter against the surplus of the moment, and the year's closing ratio can end above the threshold on renewals alone.[29]

The endowment (§A.5) is set so that the opening book's expected written premium W and the steady-state reserves of a book of that premium, $Q = (1 - \pi)\,\tau\,W$ with τ the mean time from booking to payment (§D.1.2), stand at the industry's aggregate RBC ratio, which gives an opening ratio of written premium to surplus of 0.338 in every seed. Surplus changes within a quarter by the premium written less expenses and less the reserve set up at the writer's expected loss, by the release of that reserve along the public reporting pattern, by the case reserves set up on the claims reported in the quarter and released at their settlement, by the payments made, and by investment income on surplus and reserves (§A.5). In a practice year no insurer fails; the restatement of the reserves at the close of the final practice year (§A.5) takes the provision for unreported claims at its true expectation, not conditioned on the claims already reported.

D The Kelly Decision and the Estimators

This appendix states the approach of Section 2 with its equations: the objective and the decision (§D.1), what is learned, from what data, and by what estimator (§D.2), the estimators in detail (§D.3), and the decision's variance term (§D.4).

D.1 The Objective and the Decision

This section states the objective and the decision that follows from it: the criterion (§D.1.1), the increment one account adds to expected growth (§D.1.2), what keeping an account is worth (§D.1.3), the decision, the chance of winning at a price weighed against the increment, with the lowest premium at which the insurer writes (§D.1.4), and the form in which the value of keeping an account enters the decision, with its anchoring (§D.1.5). Each quantity is stated with the data it uses, and Appendix E gives the computation in the order in which it is carried out.

D.1.1 The Objective

Let S be the insurer's surplus. The Kelly criterion is the objective: the insurer chooses its actions to maximize the expected logarithm of its surplus at the horizon, $\mathbb{E}[\ln S_T]$. Since $\ln S_T - \ln S_0$ is T times the mean yearly growth of the logarithm of surplus, the mean geometric rate of growth in logarithms, this is the maximization of that expected rate; the rate a year the tables state is the exponential, less one, of the mean of g/T, that rate net of the return on funds (Equation A.3). Where an action leaves the insurer's later opportunities unchanged, the objective is the maximization of $\mathbb{E}[\ln S']$, the expected logarithm of next period's surplus, in each period; where it changes them,

[29] National Association of Insurance Commissioners, "NAIC IRIS Ratios," data year 2025, property and casualty (NAIC, May 19, 2026).

as when an account written today is held in later years (§D.1.3), the one-period objective carries what the action leaves for later. Two features of an insurer's position change the statement of the criterion's property in repeated independent bets on known odds (§2.1). The insurer does not choose a stake at given odds: it quotes a price, the buyer decides, and the account is written whole or not at all. And the insurer does not know the distribution of an account's loss; it has an estimate. The decision below approximates the criterion for the distribution the insurer is given: to second order in the account's share of surplus, without the square of the account's expected margin, one account and one year at a time, with the value of later years entered to first order (§D.1.3). The property of optimal growth is not claimed for this approximation in this market.

D.1.2 The Increment of One Account

Consider one account with loss X, estimated mean $\hat{x} = \mathbb{E}[X \mid \text{data}]$, and estimated variance $\hat{v}$, quoted at premium P. Written, it changes surplus by $n - X$, where $n = P(1 - \varepsilon)$ is the premium net of expenses. Expanding the logarithm to second order in the account's share of surplus, the account taken alone against the current surplus, and dropping the square of the expected margin,

$$\mathbb{E}\Big[\ln \frac{S + n - X}{S}\Big] \;\approx\; \frac{n - \hat{x}}{S} - \frac{\hat{v}}{2S^2}, \tag{D.1}$$

which is a mean–variance rule in which the coefficient of the variance, one half over the surplus, is fixed by the objective and is not a parameter. The square of the expected margin is a term of the same order; the paper does not measure its size against the variance term. Two further terms enter the increment. The reserve held against the policy's expected loss earns the return on funds until its claims are paid: with τ = 4.083 years the mean time from booking to payment under the market's reporting and settlement patterns, the expected loss net of that return is $\hat{x}(1 - c_x)$ with $c_x = r\,\tau$ = 0.1454, the return over the mean time to payment without compounding. And in a market where an account written today is likely to be held in later years, the increment carries what keeping the account is worth, V per unit of net premium, the lifetime value of §D.1.3. The increment with the quote-anchored lifetime value (§D.1.5; Equation 1 gives the anchored form of the main construction) is

$$G(P) \;=\; \frac{n\,(1 + V) - \hat{x}\,(1 - c_x)}{S} - \frac{\hat{v}}{2S^2}. \tag{D.2}$$

In the main market, at the endowment's mean over the seeds, $556.7 million, the variance term is equivalent to a charge of 0.0641 percent of the expected loss of an account at the bureau's scale of claim sizes, and since the expected loss and the variance both rise in proportion to the claim rate, it does not vary with the account's size. §D.4 measures what removing it does.

D.1.3 What Keeping an Account Is Worth

Equation D.1 values an account by one year's margin. In a market where the buyer bears a cost of switching and the holder knows the account better than a challenger does, an account written today is likely to be held next year, and the objective, which is growth through the horizon, counts the margins of the years to come. The objective also fixes how they are discounted. A margin that an account adds k years ahead raises $\ln S_T$ by about that margin divided by S_{t+k}, and where the insurer's surplus grows at its own yearly rate $\hat{g}$, gross of the return on funds, $S_{t+k} \approx S_t\, e^{\hat{g}k}$, the margin counts about $e^{-\hat{g}k}$ in units of today's surplus: the discount factor is $\beta = e^{-\hat{g}}$. The step

replaces the expectation of S_t/S_{t+k} by its value on the path of mean growth, which by Jensen's inequality understates it; the funds discount of Table 1 takes another first-order route: it takes the surplus to grow at the return on funds alone, $S_{t+k} \approx S_t(1+r)^k$, so that $\beta = 1/(1+r)$, and so allows for no growth from writing. A margin that falls after the final measured year adds nothing to $\ln S_T$, so the sum runs to the horizon, over H years, the years remaining from the year after the one quoted through the final measured year. With p_1 and m_1 the retention and the expected margin per unit of net premium at an account's first renewal, and p_s and m_s at later renewals,

$$V_{\text{ren}}(H) = m_s \sum_{k=1}^{H} (\beta\, p_s)^k \qquad \text{and} \qquad V_{\text{nb}}(H) = \beta\, p_1 \big[m_1 + V_{\text{ren}}(H-1) \big], \tag{D.3}$$

with $V_{\text{nb}}(0) = 0$. This is the *lifetime value*: for an account the insurer holds, the discounted margins of the renewals to come; for an account it does not, the first renewal's margin and what follows it, reached with probability p_1. The five quantities are estimated on the insurer's own book over the lifetime value's estimation window (three years): the retentions in accounts, the margins per unit of net premium, each the numerator of Equation D.2 without the lifetime value, on first renewals kept and on later renewals kept apart—the margin of a newly won account's first renewal is the margin of a selected account, which is what the decision needs—and $\hat{g}$ as the window's mean yearly log growth of surplus, the return on funds included, each year's growth taken from the surplus at its opening to that at its close and the window taking the practice years as it takes the measured years. The derivation carries no cap: a negative margin gives a negative lifetime value, and a finite horizon, here the final measured year, needs no bound for any $\hat{g}$ (a going-concern sum exists only where $\beta p_s < 1$). Until the window holds a first renewal, the lifetime value on new business is zero.

D.1.4 The Chance of Winning, and the Decision

A quote is a bid, and the increment $G(P)$ is earned only if the quote holds the account. Let $h(\ell)$ be the probability that a quote at log deviation ℓ from the undeviated price holds (on new business, among accounts that go to market; §D.2.2). The expected increment of quoting is $h(\ell)\, G(P_\ell)$ with P_ℓ the undeviated price times e^ℓ, and the Kelly decision is the premium of the deviation that maximizes it, with G of Equation D.2 in the quote-anchored form and of Equation 1 in the anchored form of the main construction:

$$\ell^* = \arg\max_{\ell}\ h(\ell)\, G\big(P_\ell\big), \qquad P = \max\Big\{ P_{\ell^*},\ \frac{\hat{x}\,(1-c_x) + \hat{v}/(2S)}{(1-\varepsilon)\,(1+V)} \Big\}. \tag{D.4}$$

A higher deviation earns more on an account that is held and holds fewer accounts; the two are weighed by the insurer's own hold curve, taken as the posterior average that §D.2.2 states. Because the Kelly route prices through a schedule plan as the others do, its deviation, the schedule modification, and, on new business, the schedule credit together stay inside the filed band, 25 percent either way in the main market and 40 percent at the wider band of Section 6, so that ℓ^* ranges over the band and is at most $\ln 1.25 = 0.223$ in the main market and 0.336 at the wider band, on an account whose schedule modification and schedule credit are one; the zero deviation is always admissible. The second element of Equation D.4 is the *floor*: the premium at which the quote-anchored increment, Equation D.2, is zero, the lowest at which writing the account adds to expected growth in that form, with the lifetime value for the account's phase, renewal or new business, inside it; the main construction applies the floor in this form (§2.1). The floor rises with the account's variance, falls as surplus grows, and falls with the value of keeping the account; and, because the hold curve

falls with the deviation, a lower increment does not lower the chosen deviation either. So of two accounts with one expected loss the one with the wider distribution is quoted no lower, and an insurer with less surplus quotes an account of the same distribution no lower. Without the lifetime value and the return on the reserve held, the floor's premium net of expenses is the expected loss plus a loading on the variance of one over twice the surplus, the premium of the variance principle at that loading; the fair premium of a discounted-cash-flow approach to rate regulation and the allocation of capital among lines of business are two other constructions of an insurer's charge for risk.[30]

Equation D.4, over the filed band around the undeviated price, is the decision. Its inputs other than the surplus S, the expense share ε, and the return on the reserve held, c_x, are $\hat{x}$ and $\hat{v}$ for each account, the lifetime values V_{nb} and V_{ren}, the hold curve h, the undeviated price, and the band. A rating plan supplies a version of $\hat{x}$, the undeviated price, and the band; it supplies none of the others. An insurer may price any of the three estimates as a plan prices, or by the decision as a deviation from that price; the first is the plan route on that estimate and the second the Kelly route (Section 5).

D.1.5 The Form of the Lifetime Value, and Its Anchoring

Equation D.2 carries the lifetime value as V per unit of the net premium n. In the decision of §D.1.4 the premium varies with the deviation under consideration, and the form in which the lifetime value follows it is a choice of construction with a measured consequence. In the *quote-anchored lifetime value*, the lifetime value enters the increment as $Vn(\ell)$, with $n(\ell)$ the net premium of the quote at deviation ℓ, in proportion to the premium of the deviation under consideration, so that the gradient of the objective carries V times the premium's own derivative: a markup earns the value of keeping the account in proportion to its size. The continuation of a won account is invariant to today's deviation to first order: next year's quote is next year's rule applied to next year's estimate, and the buyer's decision to shop compares the holder's renewal quote with the lowest invited quote (§A.4); this year's markup enters neither quote. The invariance is argued from the market's rules; the paper does not measure each won account's continuation by the deviation that won it. What today's deviation buys is today's margin and the chance of winning the account at all: the chance of winning $h(\ell)$ of Equation D.4 multiplies the whole increment, so the value of keeping the account, $h(\ell)\,V\,\bar{n}$, is carried in full. The main construction therefore anchors the lifetime value at the undeviated price, the price of the zero deviation. This is the *anchored lifetime value*, $V\bar{n}$ in place of $Vn(\ell)$, $\bar{n} = n(0)$ the net premium of the zero deviation. It removes the gradient of the quote-anchored form, and at a deviation other than zero it also changes the lifetime value's level. The paper reports both forms, and §5.1 measures what each is worth; Appendix F gives the order in which they were registered.

The anchored form has a stated limit. The policy's own renewals are written above the zero deviation on average, so a lifetime value anchored there understates the lifetime value's level by the multiplier of the policy's mean renewal deviation, and since the chance of winning multiplies the whole increment, an understated level biases the chosen deviations upward relative to a lifetime value at the realized scale. An anchor at the realized scale would make the lifetime value a function

[30] Hans Bühlmann, *Mathematical Methods in Risk Theory* (Berlin: Springer-Verlag, 1970), chap. 4; Stewart C. Myers and Richard A. Cohn, "A Discounted Cash Flow Approach to Property-Liability Insurance Rate Regulation," in *Fair Rate of Return in Property-Liability Insurance*, ed. J. David Cummins and Scott E. Harrington (Boston: Kluwer-Nijhoff, 1987), 55–78; J. David Cummins, "Allocation of Capital in the Insurance Industry," *Risk Management and Insurance Review* 3, no. 1 (2000): 7–27.

of the deviation it prices, the self-reference the anchor removes, so the understatement is a limit of the construction: at any quote, other inputs held, the deviation chosen is higher than a lifetime value at the realized scale would give; the paper does not run a lifetime value at the realized scale, and does not measure by how much.

D.2 Learning

This section states what the two routes learn and how: each account's loss distribution, tracked from the files the insurer reads (§D.2.1), which both routes price; and the hold curve, from the broker's figures (§D.2.2), which enters the Kelly route's decision with the lifetime value of §D.1.3, whose five quantities are estimated on the insurer's own book (Appendix E). Every estimator uses only data that an insurer in the market of Section 3 holds, and what the market treats as public: the bureau's curve of claim sizes and its moments at scale one, the reporting and settlement patterns (§A.1, §A.6), the spread of the case-reserve error on a broker's loss run and on a holder's own file (§A.2, §A.6), the share of loss adjustment expense (§A.1), and, as for every insurer, the distribution of the traits within a class that the schedule modification uses (Appendix B).

D.2.1 The Account's Loss Distribution

A rating plan prices an account's history through one number, the experience modification on its capped losses. The learning tracks two quantities for each account, its claim frequency relative to its class and the scale of its claim sizes, and it does so by the best linear predictor of credibility theory, at moments it estimates, in a form that allows for a risk that changes over time.[31] For each year of a file it forms two observations: the reported claim count over its expected number at the class's frequency, and the mean logarithm of the reported claim sizes less the bureau's curve's mean log size, an open claim's case reserve corrected for its error. Its estimate of the coming year's frequency relativity and log scale is the best linear predictor from these observations jointly, in which the weight on a year depends on how much the year's observation carries (the expected claims behind it) and on how far back it lies.

The predictor needs the second moments of the two traits: how much of an account's departure from its class stays and how much decays, and at what rate. The insurer estimates them. From every file and loss run it has read it keeps the autocovariance of each observation at lags of 1 through 20 years, fits to each a constant plus a geometrically decaying term, a part that stays and a part that decays with a persistence, and keeps one covariance between the two traits (§D.3). On an account it holds it reads its own file, every year it has held the account and the loss run from before; on an account it is asked to quote it reads the broker's loss run. The expected loss is then the class's claim frequency times the account's size, times the tracked relativity, times the expected limited claim at the tracked scale, with loss adjustment expense, and the variance is that of a compound Poisson sum at the same rate and scale, which leaves out the variance from the uncertainty of the rate and the part of the scale's uncertainty common to the account's claims:

$$\hat{x}_0 = \hat{\lambda}\,\mathbb{E}\big[\min(Y,u)\big]\,(1+a), \qquad \hat{v}_0 = \hat{\lambda}\,\mathbb{E}\big[\min(Y,u)^2\big]\,(1+a)^2, \tag{D.5}$$

with $\hat{\lambda}$ the tracked claim rate, Y the claim size at the tracked scale, u the occurrence limit, and a the loss adjustment expense's share. The underwriter's schedule modification multiplies $\hat{x}_0$ as it

[31] The predictor, for a risk that does not change over time: Hans Bühlmann, "Experience Rating and Credibility," *ASTIN Bulletin* 4, no. 3 (1967): 199–207. A risk that changes over time: Howard C. Mahler, "An Example of Credibility and Shifting Risk Parameters," *Proceedings of the Casualty Actuarial Society* 77 (1990): 225–308.

does the rating plan's estimate, and $\hat{v}_0$ by the same multiplier, as a change of the claim rate would; from here on $\hat{x}_0$ and $\hat{v}_0$ denote the two after it. On either route the price carries the insurer's rate level R, so that the expected loss in the price is $\hat{x} = R\,\hat{x}_0$, and on the Kelly route the variance is $\hat{v} = R^2\,\hat{v}_0$, the rate level taken as a change of the loss's scale. The schedule modification and the tracked estimate both depend on the account's permanent traits, the first through the underwriter's reading and the second through the account's claims; their product is not adjusted for what the two share, and the rate level, moved each year by the loss-ratio method (§A.3), is the one correction of its level on the business written. The variance is the quantity a plan does not hold; it enters the Kelly route's price through the increment, Equation 1 in the main construction and Equation D.2 in the quote-anchored form, which Equation D.4 multiplies by the chance of winning and maximizes, and through the floor of Equation D.4. The Kelly route on the rating plan's estimate takes the variance in the same form, at the bureau's claim frequency for the account's class and size times the account's experience modification, with the class's claim sizes on the bureau's curve.

D.2.2 The Hold Curve

The hold curve h of Equation D.4 is learned from the only data on competitors' prices an insurer has, the broker's figure. The figure is censored: it is given on the quotes that did not bind on an account that went to market and withheld on those that bound. A challenger's quote on an account that stayed with its holder without going to market is neither: it enters no fit. On new business the hold curve is therefore the chance of winning an account that goes to market, and the decision does not weigh the chance that the account goes to market, which the challenger's own quote moves (§A.4). Let ζ be a latent log price: the logarithm of the premium at which the account would bind with another insurer, over the undeviated price for it at the quote, in each *segment*, a size band in renewals or in new business. On a quote that was lost at market, ζ is observed, the quote's own deviation less the figure. Every quote that bound enters as the event $\zeta \geq \ell$, since its holder is told nothing, including a renewal kept against a lower competing quote and one kept without going to market, which the market may also have undercut; the buyer's costs of shopping and switching are thus absorbed in the fitted μ_ζ, and h is a fitted hold curve rather than the distribution of where the market stands. The insurer models ζ as normal, $\zeta \sim N(\mu_\zeta, \sigma^2)$, with a stated prior: $\mu_\zeta \sim N(0, 0.25^2)$, the market centered on the undeviated price with the spread of the filed band, 25 percent in the main market, and $\ln\sigma \sim N(\ln 0.10, 1.0^2)$. At the wider band of Section 6 the prior on μ_ζ is $N(0, 0.40^2)$. The posterior is taken at the mode of the decayed log likelihood plus the log prior, the lost quotes entering by their density and the bound quotes by the probability of the event, with weights that decline by 2 percent a quarter, and its covariance from the inverse of the Hessian of the negative log posterior at the mode, a Laplace approximation.

The decision then uses the chance of winning averaged over that posterior,

$$\bar{h}(\ell) = \int \Big[1 - \Phi\Big(\frac{\ell - \mu_\zeta}{\sigma}\Big)\Big]\, \pi(\mu_\zeta, \sigma \mid \text{data})\, d(\mu_\zeta, \sigma), \tag{D.6}$$

exactly over μ_ζ given $\ln\sigma$ and by a 12-node Gauss–Hermite rule over $\ln\sigma$. For one period the average is what the objective implies under uncertainty about the hold curve's parameters: the increment G, in either form, does not depend on them, so the expected increment of a deviation is the posterior-mean chance of winning times G. The decision does not value the information a deviation adds about the hold curve in later quarters, and its deviations carry no rule that keeps them dispersed. In den Boer and Zwart's setting such a rule is an interval around the average of past prices, shrinking over time, that the price chosen on the current estimate may not enter, and under

it the optimal price is eventually learned.[32] Before any observation the prior predictive governs, so a quote can leave the undeviated price from the first quarter. The curve is fitted separately in each size band for renewals and for new business, six segments in all.

D.3 The Estimators

The estimator's observations are formed account-year by account-year. For an account-year s in a file let n_s be the reported claims with a positive value, q_s the expected reported claims at the class's frequency and the account's size, $\bar{y}_s$ the mean logarithm of the reported claim sizes, with $\sigma_c^2/2$ added for each open claim, and m_0 and v_0 the mean and variance of the logarithm of a claim on the bureau's curve at scale one before the occurrence limit, while the reported sizes are capped at it. The frequency observation is n_s/q_s, with noise variance $1/q_s$ around the year's relativity; the severity observation is $\bar{y}_s - m_0$ less the class's mean log scale, with noise variance $(n_s v_0 + n_s^{\text{open}}\sigma_c^2)/n_s^2$, n_s^{open} the open claims among them. A year without a reported claim carries no severity observation. The class's mean log scale is the average of $\bar{y}_s - m_0$, weighted by n_s, over the account-years of the class in every block the insurer has read, an account-year counting at each read.

Over every block of consecutive years the insurer has read, it accumulates the products of the frequency observations less one at each lag k from 1 through 20, weighted by the expected claims, and the products of the centered severity observations at each lag, a year's observation weighted by $n_s/(n_s+5)$, and it accumulates the product of the two kinds within an account. To the frequency autocovariances it fits $A_f + B_f\,\rho_f^k$ by weighted least squares over a grid of ρ_f from 0.30 through 0.98 with A_f and B_f not below zero, and likewise $A_s + B_s\,\rho_s^k$ to the severity autocovariances. The cross moment is one constant, bounded by $\sqrt{A_f A_s}$ in absolute value. It refits after every quarter's reads.

For the predictor, with an account's file of w years let o be the vector of its observations, at most $2w$, centered, Σ their covariance under the fitted moments (the fitted autocovariances between years and the cross moment between kinds, with each observation's noise variance added to the fitted variance on the diagonal), and c_f and c_s the covariances of the observations with the coming year's frequency relativity and log scale. The tracked relativity is $1 + c_f'\Sigma^{-1}o$, not below 0.05, and the tracked log scale is the class's mean plus $c_s'\Sigma^{-1}o$, plus the covariance of the frequency relativity and the log scale that remains given the observations, with variance $A_s + B_s - c_s'\Sigma^{-1}c_s$. The expected limited claim and its second moment in Equation D.5 are integrals of the bureau's curve's closed forms over a normal distribution of the log scale with that mean and variance, by Gauss–Hermite quadrature.

In the hold curve, for a segment, renewals or new business in a band, let quote j have log deviation ℓ_j from the undeviated price and weight ω_j. The log likelihood is

$$\sum_{j\,\text{lost}} \omega_j \ln\phi\Big(\frac{\zeta_j - \mu_\zeta}{\sigma}\Big) - \sum_{j\,\text{lost}} \omega_j \ln\sigma + \sum_{j\,\text{bound}} \omega_j \ln\Big[1 - \Phi\Big(\frac{\ell_j - \mu_\zeta}{\sigma}\Big)\Big], \tag{D.7}$$

with ζ_j the deviation ℓ_j less the broker's figure in logarithms and ϕ the standard normal density. With the log prior of §D.2.2 added, it is maximized in $(\mu_\zeta, \ln\sigma)$ by Newton's method, the covariance

[32] Arnoud V. den Boer, "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions," *Surveys in Operations Research and Management Science* 20, no. 1 (2015): 1–18; Arnoud V. den Boer and Bert Zwart, "Simultaneously Learning and Optimizing Using Controlled Variance Pricing," *Management Science* 60, no. 3 (2014): 770–783.

taken as §D.2.2 states. The pessimism variant of Table 1 lowers the conditional posterior mean of μ_ζ by one posterior standard deviation before the average of Equation D.6 is taken, which lowers the chance of winning at every deviation.

D.4 The Variance Term

Without the variance term the Kelly route on the learned estimate, in its main construction, chooses another deviation on 0.028 of the quotes it computes over the measured years, and the 90 percent interval of the mean of its growth less its growth with the term, paired by seed over 200 seeds, runs from -0.0059 to +0.0051 ($t = -0.13$), inside $(-0.010, +0.010)$, the range of the reading rule below, narrower than the equivalence margin of §5.1. The fifth registration (Appendix F) fixed a reading rule for this contrast before it was read: removing the term is shown not to move growth by more than 0.010 where the 90 percent interval lies inside that range; otherwise, where the interval excludes zero, the effect is reported at its sign; otherwise it is not shown. By the reading rule, removing the term is shown not to move the Kelly route's mean growth by more than 0.010. In the market of one tenth of the accounts the term is equivalent to a charge of 0.6421 percent of the expected loss of an account at the bureau's scale of claim sizes, 10.0 times its size in the main market. There the twelfth registration (Appendix F) ran the Kelly route on the learned estimate without it, read by the same rule: removing the term lowers the Kelly route's growth by 2.3 percent of terminal surplus ($t = -2.08$), and the 90 percent interval, from -0.0415 to -0.0048, excludes zero, so the effect is reported at its sign; it is not at the level that allows for the number of tests (§7.2).

E The Kelly Route's Computation, Step by Step

The steps are in the order in which they are carried out. Each names the data it uses; all of it is data the insurer holds or that the market treats as public (§D.2). The plan route, on the learned estimate and with learning on new business, carries out the first three steps of a quote and quotes the undeviated price of the third step; it also carries out the first step after each quarter's placements and the first and third steps at the year's close.

At each quote, for each account:

1. Read the file. On an account the insurer holds: its own claim file for every year held, and the loss run received at acquisition for the five years before. On any other account: the broker's loss run of five years. Read the class and the bureau's claim frequency for it.

2. Track. Form the observations of §D.3 for each year of the file and compute the relativity, the log scale, and its variance from the fitted moments. Compute $\hat{x}_0$ and $\hat{v}_0$ by Equation D.5, and multiply both by the schedule modification. On the rating plan's estimate $\hat{x}_0$ is instead the plan's expected loss of Equation A.1, its schedule modification included, without the rate level, and $\hat{v}_0$ is formed as the final sentence of §D.2.1 states. On renewals with learning on new business $\hat{x}_0$ is the plan's expected loss in the same way, and $\hat{v}_0$ is the learned variance computed above, the schedule modification included.

3. Price through the plan. Multiply $\hat{x}_0$ by the insurer's rate level R, and $\hat{v}_0$ by R^2, giving $\hat{x}$ and $\hat{v}$. The undeviated price is $R\,\hat{x}_0/(1-\pi)$, on new business times the schedule credit of the

account's size band, the schedule credit and the schedule modification together held within $1 \pm b$, b the filed band, 0.25 in the main market and 0.40 at the wider band.

4. Deviate. Take the deviation ℓ, on a grid over the filed band, and at zero, that maximizes $\bar{h}(\ell)\, G(P_\ell)$, with P_ℓ the undeviated price times e^{ℓ}, G of Equation 1 with the lifetime value for the account's phase, renewal or new business, anchored at the undeviated price, and $\bar{h}$ the posterior-averaged chance of winning of the account's segment (Equation D.6), in the main construction with the conditional posterior mean of μ_ζ lowered by one posterior standard deviation before the average, the pessimism variant (§D.3). Hold the deviation plus the logarithms of the schedule modification and the schedule credit within $\ln(1 \pm b)$; the zero deviation is always admissible.

5. Apply the floor. Quote the larger of the deviated price and the floor of Equation D.4, with S the insurer's surplus at the quote.

After each quarter's placements:

1. Add every file and loss run read in the quarter to the sums of the second moments, and refit them.

2. For the segment's hold curve, record each quote that bound with its deviation, among them a renewal kept without going to market and one kept against a lower competing quote, and each quote lost on an account that went to market with its deviation and the broker's figure; decay the weights by 0.98 a quarter and refit the posterior of each segment.

3. For each renewal quoted, record whether it was a first renewal and whether it was kept; for each kept, the premium net of expenses and Equation D.2's margin, first renewals and later ones apart.

At each year's close:

1. Move the rate level by the loss-ratio method on the insurer's own experience, with competitors' winning prices as the complement of credibility (Equation A.2, §A.3).

2. Update the lifetime value. Over the past three years, take p_1 and p_s as the shares of first and of later renewals kept, and m_1 and m_s as the margins per unit of net premium on each; set β at $1/(1+r)$ on the funds discount, or at $e^{-\hat{g}}$ on the growth discount, $\hat{g}$ the window's mean yearly log growth of surplus, the return on funds included; and set H of §D.1.3, the years remaining to the horizon, and V_{ren} and V_{nb} by Equation D.3.

3. Update the schedule credit of each size band from the broker's figures on the insurer's own quotes on new business (Appendix B).

F The Registry of Predictions and Outcomes

Every prediction of the paper was written down before the runs that tested it, or, for the sixth registration, which analyzes runs already made, before any of its metrics was computed, under one rule: a predicted contrast holds where its paired mean over the 200 seeds has the predicted sign at $|t| > 2$, and a prediction that names another reading is read as it names. This appendix states the

12 registrations in the order of their writing, each with its predictions on the insurers this paper reports, their outcomes under the rule, the order of writing against the order of running, and any reading of other seeds that preceded it; the body points here for the outcomes it reports. The registrations also carried 101 predictions on an insurer that prices by the decision alone, with no rate level, on its own learned estimate and on the plan's, on the market around it, and on a program of variations of the market run on it, which this paper does not report: 91 of the first registration's, three of the second's, all four of the fourth's, and three of the fifth's. Of them 82 held, 18 failed, and one, the first registration's, was not tested.

The first registration made 97 predictions, registered after a reading of 20 other seeds and before any arm ran on the 200 seeds. The six reported here held: the plan route on the learned estimate against the benchmark in the main market and at each of the four other settings of Section 4 (Table F.1), and, in the market with the benchmark, the buyer's calibration: on the 200 seeds 0.141 of that market's accounts change insurer a year (standard error over seeds 0.0009), within two standard errors of the target of 0.14 (§A.4).

Table F.1: The plan route's contrast at the five settings of Section 4 and its leads over the conventional insurers, paired by seed over 200 seeds: the mean on the logarithm (Equation A.3), the same as a difference in terminal surplus in percent (the exponential less one), its t statistic, and the seeds above zero. The first contrast is the plan route on the learned estimate less the benchmark, the fourth seat's growth in two arms run on the same seeds; its block's first row is the main market, and each of the next four rows moves one of two of the main market's dimensions. The second, which no registration predicts, is the plan route on the learned estimate less each conventional insurer in the main market. Contrasts are taken on unrounded values and may differ by 0.001 from differences of the means Figure 2 states. Source: authors' calculations.

Contrast	Mean	Percent	t	Seeds above zero
Learning the loss: the plan route on the learned estimate less the benchmark				
The main market	$+0.400$	$+49.2$	$+31.10$	197
Refiling threshold 0.05 (lower)	$+0.394$	$+48.3$	$+31.42$	198
Refiling threshold 0.15 (higher)	$+0.391$	$+47.9$	$+31.57$	195
Read noise 0.35 (lower)	$+0.490$	$+63.3$	$+36.08$	199
Read noise 1.05 (higher)	$+0.384$	$+46.8$	$+32.36$	196
The plan route on the learned estimate less a conventional insurer				
Bureau insurer	$+0.541$	$+71.8$	$+42.09$	200
Loss-ratio insurer	$+0.612$	$+84.4$	$+48.62$	198
GLM insurer	$+0.519$	$+68.0$	$+43.39$	200

The second registration made 10 predictions, written after the first registration's outcomes were known. The seven reported here were written after a reading of other seeds: one, that its arms carried over from the first registration, the plan route's among them, reproduce the first registration's runs, run for run; and six on four directions first read on those seeds, the anchoring on the Kelly route, its three leads over the conventional insurers, the pessimism variant on it, and the Kelly route with pessimism against the plan route, the primary prediction. The reading ran six constructions of the Kelly route. Two had the quote-anchored lifetime value, one on each discount of the lifetime value. Four changed it on the funds discount: without the lifetime value; with the anchored lifetime value; with the pessimism variant; and with both the anchored lifetime value and the pessimism variant. Of the seven, five held and two failed: the pessimism variant on the Kelly route and the primary prediction.

The third registration made two predictions for the Kelly route on the growth discount, written before any run of the arms it tests, and both failed: the growth discount's effect on the Kelly route without pessimism and the primary prediction evaluated on that discount. The fourth registration's four predictions are among those the count above names.

The fifth registration ran the four further markets beside the main market and made six predictions for two of them, the market of smaller insurers and the market with a second insurer that learns, written after a reading of 20 other seeds and before the arms ran on the 200 seeds. The three reported here held: the plan route on the learned estimate against the benchmark in each of the two markets, and, in the market with a second insurer that learns, that insurer's leads where the fourth seat is the benchmark (Appendix G). It made no prediction for the Kelly route, which it ran in the market with a second insurer that learns and under the severity trend, and for the Kelly route without the variance term it states a reading rule in place of predictions (§D.4).

The sixth registration is an analysis on the records in hand, with no run of the market: the equivalence margin of §5.1 and the metrics of the two routes' paths on the learned estimate (§5.1; §5.2 reports the same metrics with learning on new business), with six predictions written before any metric was computed. The primary contrast's value was known and is stated in the registration as known, and the difference beside a second insurer that learns was on the record from the fifth registration; the margin is 0.028 in the logarithm over the 28 measured years, 2.8 percent of terminal surplus and a tenth of a percentage point of the mean geometric rate of growth a year, chosen in the objective's own unit and set below the smallest effects on which the registered findings made before it rest, the anchoring's worth on the insurer that prices by the decision alone, which the paper does not report: 0.054 in the logarithm on the growth discount and 0.033 on the funds discount; a margin of a tenth of learning's worth would have placed the second inside it. Registered findings made after it rest on smaller effects, among them the seventh registration's sixth prediction and the eleventh's primary prediction at the read noise of 0.35. The first prediction is the equivalence reading itself at that margin; the second and third are read at $t > 2$, the paired within-path volatility and the Pitman–Morgan t of the dispersion across seeds; the fourth through sixth are read on medians and shares of the paired path, without a t. The equivalence was shown; of the five predictions that name a direction, four held and one failed, that on the within-path volatility, which the Kelly route has below the plan route's where above was predicted.

The seventh registration ran the two routes on the learned estimate at other values of the filed band and the buyer's scale (Section 6), with six predictions written before its reading of 20 other seeds, which left them unchanged, and before any arm ran on the 200 seeds: the first, that its runs reproduce the reported arms at the main market's values run for run, read on 60 runs on other seeds and held there; the second, a loss for the Kelly route against the plan route at the wider band; the third and fourth, that the buyer's scale moves the Kelly route's mean renewal deviation and its band-edge share below their values at the main market's scale at the lower scale and above them at the higher; and the fifth and sixth, that the routes' contrast less the same contrast at the main market's scale is below zero at the lower scale and above it at the higher. All six held on the 200 seeds; at the reading of 20 other seeds one of them, the sixth, had failed. The registration's Kelly-route arms at the wider band kept the hold curve's grid of deviations and prior at the 25 percent band, and its second prediction was read on the arm of the main construction; Table 6 gives in that arm's place the twelfth registration's run of the same construction, whose hold curve is set at the wider band (below); it also trails. The seventh registration also ran the Kelly route in the market of smaller insurers and the market of one tenth of the accounts, without a prediction.

Table F.2: The contrasts among the insurers of Table 1 on the learned estimate that the second and third registrations read, the fourth seat's growth paired by seed over 200 seeds: the mean on the logarithm (Equation A.3), its t statistic, and the seeds above zero. The anchoring is an arm with the anchored lifetime value less the same arm with the quote-anchored lifetime value, the pessimism variant's contrast an arm with the variant less the same arm without it, and the growth discount's contrast an arm on the growth discount less the same arm on the funds discount. Above the subheading *On the growth discount* are the second registration's contrasts, on the funds discount: its predictions are on the anchoring, the pessimism variant, and the Kelly route with pessimism less the plan route, and the block's two other rows carry none. Below it are the third registration's, on the growth discount: its two predictions are on the growth discount's contrast on the Kelly route without pessimism and on the Kelly route with pessimism less the plan route. The four contrasts with the plan route that Table F.3 also gives are printed at its four decimals. Under the registered rule a predicted contrast holds where its paired mean has the predicted sign at $|t| > 2$; contrasts are taken on unrounded values. Source: authors' calculations.

Contrast, the fourth seat's growth	Mean	t	Seeds above zero
The anchoring, on the Kelly route without pessimism	$+0.110$	$+14.85$	177
Pessimism, on the Kelly route	0.000	$+0.03$	100
Less the plan route			
The Kelly route with pessimism	$+0.0108$	$+1.97$	113
The Kelly route without pessimism	$+0.0107$	$+1.92$	118
The Kelly route without pessimism, quote-anchored lifetime value	-0.099	-12.01	34
On the growth discount			
The growth discount, on the Kelly route without pessimism	-0.006	-1.36	86
The growth discount, on the same with pessimism	-0.005	-0.96	99
Pessimism, on the Kelly route	$+0.002$	$+0.42$	98
Less the plan route			
The Kelly route with pessimism	$+0.0062$	$+1.03$	119
The Kelly route without pessimism	$+0.0043$	$+0.71$	113

The eighth registration ran the Kelly route on the rating plan's estimate, the cell of the grid not run before it, with four predictions written before a reading of 20 other seeds and before any arm ran on the 200 seeds. The first, that its runs reproduce the reported arms at the main market's values run for run, was read on the runs on other seeds and held there. The second, the primary, was that the new cell's growth is equivalent at the margin to the benchmark's. It failed: the Kelly route on the rating plan's estimate grows more, by 21.8 percent of terminal surplus (Section 5). The third, that learning the loss adds to the Kelly route's growth, held (23.8 percent, t = +19.44). The fourth, that the Kelly route on the rating plan's estimate quotes renewals above and new business below its undeviated price, held: its mean renewal deviation is 2.6 percent above that price and its mean new-business deviation 10.8 percent below.

The ninth registration decomposed learning, with eight predictions written before a reading of 20 other seeds and before any arm ran on the 200 seeds (Section 4, Table 2). The first, that its runs reproduce the reported arms, held on other seeds. The second, the primary, that learning the claim frequency is worth more than learning the claim size, held, and so did the third, that each is worth something learned alone. The fourth, that the benchmark, pricing on the rating plan's estimate, grows more than the fourth seat at the class value of both traits, held, as did the eighth, the experience modification's worth over the class loss cost. The fifth predicted that learning the loss adds both on new business alone and on renewals alone. It failed: on renewals alone learning

the loss takes away 15.1 percent of terminal surplus. The sixth, that learning on new business alone is worth more than on renewals alone, held. The seventh, that the true moments are worth more than the estimated, held on the 200 seeds; at the reading of other seeds it had failed.

The tenth registration ran the Kelly route with learning on new business, with three predictions. Its reading of 20 other seeds was launched before the predictions were written and read after. The first, that its runs reproduce the reported arms, held on other seeds. The second, the primary, that learning on new business adds to the Kelly route's growth, failed with the opposite sign: it takes away 2.7 percent of terminal surplus ($t = -3.52$, not at the level that allows for the number of tests, §7.2). The third, that with learning on new business the two routes end the measured years with equivalent growth, failed: the Kelly route trails by 9.6 percent (Section 5).

The eleventh registration ran the plan route with learning on new business at the other settings and in the four further markets of Section 6, the Kelly route with learning on new business at three of the other settings and in the four further markets, and the Kelly route's four variations with learning on new business in the main market, with five predictions written before any arm ran. The first, that its runs at seed 0 reproduce the reported arms they are compared with, held. The second through fourth were read at each setting and in each further market they name, and the fifth for each variation. The second, the primary, was that learning on new business adds to the plan route on the learned estimate at each of the seven other settings and the four further markets. It held at eight of the eleven. It failed beside a second insurer that learns and under the severity trend, where the difference has the opposite sign (under the trend, not at the level that allows for the number of tests, §7.2), and in the market of one tenth of the accounts, where it is not resolved (Table 5). The third, that the plan route with learning on new business grows more than the benchmark at each of the four other settings and the four further markets where the benchmark was run, held at all eight. The fourth, that with learning on new business the Kelly route trails the plan route at each of the three other settings and the four further markets where both routes were run, held at six of the seven; in the market of one tenth of the accounts the difference is not resolved (Table 6). At the wider band the fourth was read on the eleventh registration's Kelly-route arm, which kept the hold curve's grid and prior at the 25 percent band; Table 6 gives in its place the twelfth registration's run, below. The fifth, that each of the Kelly route's four variations trails the plan route with learning on new business, held for all four (§5.1).

The twelfth registration ran the Kelly route at the wider band again, with its hold curve set at that band, on the learned estimate and with learning on new business. It also ran the Kelly route on the learned estimate without the variance term in the market of one tenth of the accounts (§D.4). Its three predictions were written after the Kelly route on the learned estimate had been run at the wider band on one seed outside the 200, in a market of reduced size with fewer practice and measured years, with its hold curve at 25 percent and at the band. In those runs the route grew more with its hold curve at the band. The predictions were written before any arm ran on the 200 seeds. The first, that its runs at seed 0 reproduce the five reported arms they are compared with, held. The second and third, that the hold curve set at the band adds growth on the learned estimate and with learning on new business, failed with the opposite sign. With its hold curve at the band the Kelly route grows less than the seventh and eleventh registrations' arms, whose hold curve was at 25 percent: by 6.4 percent of terminal surplus on the learned estimate, against the seventh's ($t = -15.76$), and by 2.0 percent with learning on new business, against the eleventh's ($t = -4.10$). The Kelly route's figures in the rows of Tables 6 and G.2 at the wider band are its runs.

The rule is applied to each prediction without adjustment for the number of predictions tested. A Bonferroni level divides the rule's two-sided 5 percent among a family of predictions read by a

t, which is every prediction except those read otherwise: the first registration's calibration and its prediction not tested; the predictions of the second, seventh, eighth, ninth, tenth, eleventh, and twelfth registrations that arms reproduce arms; the sixth's equivalence reading and its three predictions on medians and shares; and the equivalence readings of the eighth and tenth. A prediction that rests on several contrasts, the leads over the three conventional insurers among them, is one prediction of its family, and it holds at an adjusted level where each of its contrasts clears it. The eleventh registration's predictions are read at each setting, in each further market, and for each variation, and each reading is one member of the family. Over the 165 members of that family across the 12 registrations, at $|t| > 3.68$, every prediction reported here that held still holds but four: the seventh registration's sixth, and three readings of the eleventh's, learning on new business against the plan route on the learned estimate at the read noise of 0.35 (t = +2.52) and, under the severity trend, learning on new business against the benchmark (t = +3.24) and, with learning on new business, the Kelly route against the plan route (t = −2.41).

Table F.3: The two routes' difference on the learned estimate at the close of the measured years and the equivalence test of the sixth registration: the Kelly route less the plan route, the fourth seat's growth paired by seed over 200 seeds, on each discount of the lifetime value with and without the pessimism variant. The mean is given on the logarithm (Equation A.3) and as the difference in terminal surplus in percent, the exponential of the mean less one; the standard error, the t statistic, and the 90 percent interval are on the logarithm, the interval the mean plus and minus 1.645 standard errors. *Within* means that the interval lies inside $(-m, +m)$ at the registered margin m, which was registered with the first row's value already measured (the sixth registration, above). Source: authors' calculations.

Less the plan route	Mean, logarithm	Percent of terminal surplus	s.e.	t	90 percent interval	At m = 0.028
The Kelly route						
With pessimism, the funds discount	+0.0108	+1.1	0.0055	+1.97	[+0.0018, +0.0198]	within
The funds discount	+0.0107	+1.1	0.0055	+1.92	[+0.0015, +0.0198]	within
With pessimism, the growth discount	+0.0062	+0.6	0.0061	+1.03	[−0.0037, +0.0162]	within
The growth discount	+0.0043	+0.4	0.0061	+0.71	[−0.0057, +0.0144]	within

G The Quote Data and the Further Markets

Table G.1 gives the quotes of the Kelly route on the learned estimate by arm and phase (§5.2 reports its mean deviations); Table G.2 gives the quotes of the Kelly route with learning on new business in the main market and at each setting and in each further market of Section 6 with both routes; and Table G.3 gives the contrast of the plan route on the learned estimate in the four further markets of the fifth registration, which the paragraph before it states.

Table G.1: The quotes of the Kelly route on the learned estimate over the measured years and the 200 seeds, by arm and phase: the mean deviation from its undeviated price, pooled over every quote of the measured years and seeds (a logarithm; §5.2 states it in percent of the undeviated price, the exponential less one); the share of quotes within one grid step (0.0025 in log price) of the filed band's upper end, the share at exactly the zero deviation, and the share raised to the floor; for the arms with the anchored lifetime value, the posterior-mean chance of winning at the chosen deviation (predicted hold); and the share of the renewal quotes that held the account (realized hold). Above the subheading, on the funds discount, are the arm with the quote-anchored lifetime value, with the anchored lifetime value, and with the anchored lifetime value and the pessimism variant, the arms of Table 1. On new business the predicted hold is over every account the arm priced and the measured share held is over the quotes it made, so no realized hold is set beside it. Below the subheading, the third registration's arms with the anchored lifetime value, on the growth discount. Source: authors' calculations.

Arm	Phase	Mean deviation	At the band's upper end	At zero	Raised to the floor	Predicted hold	Realized hold
Quote-anchored	Renewals	+0.174	0.288	0.127	0.000	—	0.753
	New business	+0.021	0.105	0.169	0.000	—	—
Anchored	Renewals	+0.120	0.278	0.166	0.000	0.775	0.809
	New business	−0.081	0.039	0.136	0.000	0.708	—
Anchored, with pessimism	Renewals	+0.118	0.276	0.168	0.000	0.776	0.811
	New business	−0.088	0.037	0.133	0.000	0.714	—
On the growth discount							
Anchored	Renewals	+0.135	0.288	0.150	0.000	0.760	0.793
	New business	−0.057	0.039	0.146	0.000	0.680	—
Anchored, with pessimism	Renewals	+0.133	0.288	0.152	0.000	0.761	0.795
	New business	−0.063	0.037	0.145	0.000	0.686	—

Table G.2: The quotes of the Kelly route with learning on new business in the main market and at each setting and in each further market of Section 6 with both routes, each seed's mean over its measured years and the mean of those over the 200 seeds (Table 4 pools the main market's quotes): the mean renewal deviation from its undeviated price (a logarithm; §5.2 states the main market's in percent of the undeviated price, the exponential less one), the share of renewal quotes within one grid step (0.0025 in log price) of the filed band's upper end, and the mean new-business deviation from that price (a logarithm). The settings and markets are those of Table 6; at the wider band the Kelly route's hold curve is set at that band (Section 6). Source: authors' calculations.

The market	Renewal deviation	Renewal quotes at the band's upper end	New-business deviation
The main market	+0.041	0.134	+0.006
Filed band of 40 percent (wider)	+0.038	0.062	+0.126
Buyer's scale 0.025 (more responsive)	+0.035	0.126	+0.061
Buyer's scale 0.075 (less responsive)	+0.042	0.135	−0.026
Smaller insurers	+0.043	0.131	+0.002
One tenth of the accounts	+0.047	0.157	+0.060
A second insurer that learns	+0.041	0.122	−0.017
Severity trend of 6.8 percent a year	+0.036	0.166	+0.184

In each of the four further markets the plan route on the learned estimate outgrows the benchmark at $|t| > 2$ (Table G.3); in the market of smaller insurers and beside a second insurer that

learns this was a registered prediction. In the market with a second insurer that learns and the benchmark, the second insurer that learns leads the Bureau insurer, the Loss-ratio insurer, and the fourth seat, as predicted (+0.505, +0.624, and +0.612 in the logarithm; t of +35.30, +44.34, and +48.19); in its arm whose fourth seat is the plan route on the learned estimate, the difference in growth between the two seats that learn is not resolved (the fourth seat less the second, -0.019 in the logarithm, $t = -0.79$). Tables 5 and 6 give learning on new business and the Kelly route against the plan route in the four further markets. Appendix F states the fifth and the eleventh registrations' predictions in the four further markets.

Table G.3: The four further markets of the fifth registration: E, every insurer at an endowment at which a seat writing the opening book stands at the capital ratio at which broker flight begins; C, one tenth of the accounts; L, the GLM insurer's seat taken by a second insurer that learns; T, a trend of 6.8 percent a year in the mean claim limited per occurrence. The plan route on the learned estimate less the benchmark, the fourth seat's growth (Equation A.3) paired by seed over 200 seeds, each mean on the logarithm above its t statistic in parentheses; Tables 5 and 6 give learning on new business and the Kelly route against the plan route in each. The variance term at each market's mean endowment is stated as a charge in percent of the expected loss of an account at the bureau's curve's own scale of claim sizes. Source: authors' calculations.

The market	E	C	L	T
The plan route on the learned estimate less the benchmark	$+0.451$ ($+26.37$)	$+0.241$ ($+12.01$)	$+0.340$ ($+23.11$)	$+0.163$ ($+7.39$)
Variance term as a charge, percent of the expected loss	0.1372	0.6421	0.0641	0.0644

Bibliography

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy* 132, no. 3 (2024): 723–771.

Association for Cooperative Operations Research and Development. *Commercial Insurance Application*. ACORD 125, 2007/10.

Bailey, Robert A. "Insurance Rates with Minimum Bias." *Proceedings of the Casualty Actuarial Society* 50 (1963): 4–11.

Breiman, Leo. "Optimal Gambling Systems for Favorable Games." In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, edited by Jerzy Neyman, 65–78. Berkeley: University of California Press, 1961.

Brown, Zach Y., and Alexander MacKay. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics* 15, no. 2 (2023): 109–156.

Bühlmann, Hans. "Experience Rating and Credibility." *ASTIN Bulletin* 4, no. 3 (1967): 199–207.

Bühlmann, Hans. *Mathematical Methods in Risk Theory*. Berlin: Springer-Verlag, 1970.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review* 110, no. 10 (2020): 3267–3297.

Capen, Edward C., Robert V. Clapp, and William M. Campbell. "Competitive Bidding in High-Risk Situations." *Journal of Petroleum Technology* 23, no. 6 (1971): 641–653.

Cummins, J. David. "Allocation of Capital in the Insurance Industry." *Risk Management and Insurance Review* 3, no. 1 (2000): 7–27.

D'Arcy, Stephen P., and Neil A. Doherty. "Adverse Selection, Private Information, and Lowballing in Insurance Markets." *Journal of Business* 63, no. 2 (1990): 145–164.

D'Arcy, Stephen P., and Neil A. Doherty. "The Aging Phenomenon and Insurance Prices." *Proceedings of the Casualty Actuarial Society* 76 (1989): 24–44.

de Garidel-Thoron, Thomas. "Welfare-Improving Asymmetric Information in Dynamic Insurance Markets." *Journal of Political Economy* 113, no. 1 (2005): 121–150.

den Boer, Arnoud V. "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions." *Surveys in Operations Research and Management Science* 20, no. 1 (2015): 1–18.

den Boer, Arnoud V., and Bert Zwart. "Simultaneously Learning and Optimizing Using Controlled Variance Pricing." *Management Science* 60, no. 3 (2014): 770–783.

Fisher, Ginda Kaplan, Lawrence McTaggart, Jill Petker, and Rebecca Pettingell. *Individual Risk Rating*. Casualty Actuarial Society study note, 2019.

Greenwich Associates. *Insurance Market Pulse: Looking at the Small Business and Middle Market Segments*. 2007. As reported in Douglas McLeod, "Changing Insurers Yields Surprise," *Business Insurance*, June 17, 2007.

The Hartford Insurance Group, Inc. *Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2025*. Filed with the U.S. Securities and Exchange Commission, February 20, 2026.

Hays, Lisa A. "Managing Commercial Lines Pricing Levels in a Loss Cost Environment." *Casualty Actuarial Society Forum* (Winter 2002): 9–30.

Hornik, Kurt, Maxwell Stinchcombe, and Halbert White. "Multilayer Feedforward Networks Are Universal Approximators." *Neural Networks* 2, no. 5 (1989): 359–366.

Jia, Ruo, and Zenan Wu. "Insurer Commitment and Dynamic Pricing Pattern." *Geneva Risk and Insurance Review* 44, no. 1 (2019): 87–135.

Keatinge, Clive L. "Modeling Losses with the Mixed Exponential Distribution." *Proceedings of the Casualty Actuarial Society* 86 (1999): 654–698.

Kelly, John L., Jr. "A New Interpretation of Information Rate." *Bell System Technical Journal* 35, no. 4 (1956): 917–926.

Kunreuther, Howard, and Mark Pauly. "Market Equilibrium with Private Knowledge: An Insurance Example." *Journal of Public Economics* 26, no. 3 (1985): 269–288.

Longley-Cook, Laurence H. "An Introduction to Credibility Theory." *Proceedings of the Casualty Actuarial Society* 49 (1962): 194–221.

Lynch, Jim, and William Nibbelin. *Increasing Inflation on Liability Insurance – Impact as of Year-End 2024*. Insurance Information Institute and Casualty Actuarial Society, October 2025.

MacLean, Leonard C., Edward O. Thorp, and William T. Ziemba, eds. *The Kelly Capital Growth Investment Criterion: Theory and Practice*. Singapore: World Scientific, 2011.

MacLean, L. C., W. T. Ziemba, and G. Blazenko. "Growth versus Security in Dynamic Investment Analysis." *Management Science* 38, no. 11 (1992): 1562–1585.

Mahler, Howard C. "An Example of Credibility and Shifting Risk Parameters." *Proceedings of the Casualty Actuarial Society* 77 (1990): 225–308.

Milgrom, Paul R., and Robert J. Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50, no. 5 (1982): 1089–1122.

Morgan, W. A. "A Test for the Significance of the Difference between the Two Variances in a Sample from a Normal Bivariate Population." *Biometrika* 31, no. 1/2 (1939): 13–19.

Myers, Stewart C., and Richard A. Cohn. "A Discounted Cash Flow Approach to Property-Liability Insurance Rate Regulation." In *Fair Rate of Return in Property-Liability Insurance*, edited by J. David Cummins and Scott E. Harrington, 55–78. Boston: Kluwer-Nijhoff, 1987.

National Association of Insurance Commissioners. "Aggregated P&C RBC Data: 2025 Data as of June 1, 2026." *Summary: Aggregate P/C RBC Results by Year*. NAIC, 2026.

National Association of Insurance Commissioners. "NAIC IRIS Ratios." Data year 2025, property and casualty. NAIC, May 19, 2026.

National Association of Insurance Commissioners. *Report on Profitability by Line by State in 2024*. Kansas City, MO: NAIC, 2026.

National Association of Insurance Commissioners. *Risk-Based Capital (RBC) for Insurers Model Act*. Model 312. 2012.

National Association of Insurance Commissioners. *U.S. Property & Casualty and Title Insurance Industries—2025 Full Year Results*. NAIC Financial Regulatory Services, 2026.

National Association of Insurance Commissioners, Casualty Actuarial and Statistical (C) Task Force. *Price Optimization White Paper*. Adopted by the task force, November 19, 2015.

Nilssen, Tore. "Consumer Lock-in with Asymmetric Information." *International Journal of Industrial Organization* 18, no. 4 (2000): 641–666.

Palmer, Joseph M. *Increased Limits Ratemaking for Liability Insurance*. Casualty Actuarial Society study note, 2006.

Pitman, E. J. G. "A Note on Normal Correlation." *Biometrika* 31, no. 1/2 (1939): 9–12.

Samuelson, Paul A. "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long." *Journal of Banking and Finance* 3, no. 4 (1979): 305–307.

Thorp, Edward O. "The Kelly Criterion in Blackjack Sports Betting, and the Stock Market." In *Handbook of Asset and Liability Management*, vol. 1, edited by Stavros A. Zenios and William T. Ziemba, 385–428. Amsterdam: North-Holland, 2006.

The Travelers Companies, Inc. "Travelers Reports Excellent Second Quarter and Year-to-Date Results." Press release, July 17, 2026.

Verisk. *2024 Executive Insights: General Liability*. 1st ed. Verisk's Core Lines Services Report, Second-Quarter 2019 to First-Quarter 2024. Verisk, 2024.

Weiss, Mary A. "Underwriting Cycles: A Synthesis and Further Directions." *Journal of Insurance Issues* 30, no. 1 (2007): 31–45.

Werner, Geoff, and Claudine Modlin. *Basic Ratemaking*. 5th ed. [Arlington, VA]: Casualty Actuarial Society, 2016.

Wu, Cheng-Sheng Peter, and Hua Lin. "Large Scale Analysis of Persistency and Renewal Discounts for Property and Casualty Insurance." *Casualty Actuarial Society E-Forum* (Winter 2009): 396–408.